



**VALUE-FOCUSED THINKING IN THE PRESENCE OF WEIGHT
AMBIGUITY: A SOLUTION TECHNIQUE USING MONTE CARLO
SIMULATION**

THESIS

Timothy W. Porter, Captain, USAF

AFIT/GOR/ENS/04-15

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

AFIT/GOR/ENS/04-15

VALUE-FOCUSED THINKING IN THE PRESENCE OF WEIGHT AMBIGUITY: A
SOLUTION TECHNIQUE USING MONTE CARLO SIMULATION

THESIS

Presented to the Faculty

Department of Operational Sciences

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the
Degree of Master of Science in Operations Research

Timothy W. Porter, B.S.

Captain, USAF

June 2004

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

VALUE-FOCUSED THINKING IN THE PRESENCE OF WEIGHT AMBIGUITY: A
SOLUTION TECHNIQUE USING MONTE CARLO SIMULATION

Timothy W. Porter, B.S.
Captain, USAF

Approved:

Stephen P. Chambal, Maj, USAF (Chairman)
Associate Professor of Operations Research
Department of Operational Sciences
Air Force Institute of Technology

date

Dr. Richard F. Deckro (Member)
Professor of Operations Research
Department of Operational Sciences
Air Force Institute of Technology

date

Acknowledgments

I would like to be able to say something nice and inspirational here. I would like to be able to say a lot of things. Some things can't be said because of tact, others can't be said out of honesty. But here are a few things that can be said...

AFIT is a great education and I thank the Air Force and the leadership at AFIT for letting me be a part of the program, and for not kicking me out. Honesty compels me to thank two men with seemingly unlimited patience, Dr. Deckro and Maj Chambal. Despite the head banging frustration and the days of utter disbelief, they didn't give up on me. At least not entirely. Thank you, I truly couldn't have done this without your support. You are both, each in your own way, excellent mentors and advisors. I would like to thank Col Mercer for taking a chance on an unknown quantity. I would also like to thank Maj Lyle and the folks at AFLMA/LGY for having a fresh perspective on some old issues.

Finally, I thank my family, who because of some bizarre rule I can't actually name. That's ok, they know who they are. Without them I would have given up a long time ago. With their support and faith, anything is possible. Thank you and I love you all.

With incredibly mixed emotions, yours always,

Tim "Never give up... EVER" Porter

Contents

	Page
<i>Acknowledgments</i> _____	<i>iv</i>
<i>List of Tables</i> _____	<i>vii</i>
<i>List of Figures</i> _____	<i>viii</i>
<i>Abstract</i> _____	<i>ix</i>
1. Introduction _____	1
1.1 Background _____	2
1.2 Problem Statement _____	5
1.3 Research Direction _____	7
1.3.1 Objectives. _____	7
1.3.2 Approach _____	8
1.3.3 Scope and Assumptions _____	9
1.4 Research Contribution _____	10
1.5 Thesis Overview _____	10
2. Literature Review _____	12
2.1 Decision Analysis _____	13
2.1.1 Utility _____	15
2.1.2 Value _____	18
2.2 Value-Focused Thinking (VFT) _____	21
2.3 Weights _____	25
2.4 Direct Weighting _____	27
2.5 Independent Scale Weighting _____	28
2.6 Swing Weighting _____	29
2.7 Decision Making in the Presence of Ambiguity _____	30
2.7.1 Uncertainty in Measures _____	31
2.7.2 Decision Makers _____	31
2.7.3 Weight Ambiguity _____	32
2.8 Monte Carlo Simulation _____	33
2.9 Selecting Alternatives in the Presence of Ambiguity _____	36
2.10 Examples _____	36
2.10.1 Mexico City Airport Citing _____	37
2.10.2 Information Assurance _____	38

3. Methodology	41
3.1 Simulation	41
3.2 Direct Weighting	43
3.3 Independent Scale Weighting	50
3.4 Swing Weighting	51
3.5 Selecting Alternatives in the Presence of Ambiguity	52
3.5.1 Dominance	53
3.5.2 Statistical Tests	54
3.6 Methodology Summary	56
4. Analysis	60
4.1 Airport Citing	60
4.1.1 Building Model	61
4.1.2 Determining n	63
4.1.3 Simulation Results	64
4.1.4 Comparing Results	67
4.2 Information Assurance	70
4.2.1 Building the Model	71
4.2.2 Determining n	73
4.2.3 Simulation Results	74
4.2.4 Dominance	78
5. Conclusions and Discussion	81
5.1 A New Approach	81
5.2 Results	82
5.3 Future Work	83
Bibliography	84

List of Tables

	Page
<i>Table 1: Attribute weight data, from Lavelle et al., 1997: 779</i> _____	62
<i>Table 2: Rating data for Alternative A, from Lavelle et al., 1997: 780</i> _____	62
<i>Table 3: Rating data for Alternative B, from Lavelle et al., 1997: 780</i> _____	62
<i>Table 4: Rating data for Alternative C, from Lavelle et al., 1997: 780</i> _____	63
<i>Table 5: Simulating Lavelle</i> _____	64
<i>Table 6: Information Assurance Simulation Results</i> _____	74

List of Figures

	Page
<i>Figure 1: Notional Value Hierarchy for a Vehicle Purchase</i> _____	22
<i>Figure 2: Weighted Information Assurance Hierarchy (Beauregard, 2001)</i> _____	39
<i>Figure 3: Notional Car Example, with weights</i> _____	45
<i>Figure 4: Direct Weighting, sum of weights</i> _____	45
<i>Figure 5: Filtered Car Value Hierarchy</i> _____	47
<i>Figure 6: Cost weight for filtered car example</i> _____	47
<i>Figure 7: Cost weight, symmetric weights</i> _____	48
<i>Figure 8: Cost Weight for Normalized Car Example</i> _____	49
<i>Figure 9: Stochastic Dominance</i> _____	54
<i>Figure 10: Methodology Flow Chart</i> _____	57
<i>Figure 11: Alternative A</i> _____	65
<i>Figure 12: Alternative B</i> _____	66
<i>Figure 13: Alternative C</i> _____	66
<i>Figure 14: IA Hierarchy with Swing Weight Coefficient Distributions</i> _____	72
<i>Figure 15: IA – Baseline</i> _____	75
<i>Figure 16: IA – ISS</i> _____	76
<i>Figure 17: IA – ESM</i> _____	76
<i>Figure 18: IA – Cisco</i> _____	77
<i>Figure 19: Information Assurance Alternatives</i> _____	78
<i>Figure 20: Cumulative distribution chart showing stochastic dominance</i> _____	80

Abstract

When a Decision Maker is asked to provide his or her preferences, the response represents a snapshot in time. While their preference structure elicited at any given moment represents their revealed preferences at that point in time, it may change over time. These changing preferences over time represent ambiguity in the decision maker's preferences. Other sources of ambiguity may result from the presence of groups in the decision making process.

One weakness of many decision analysis techniques today is the inability to incorporate this possible ambiguity into the basic decision model. The existence of the problem has been known and commented on for many years. This research attempts to address that problem. It begins with the basic approach and methodology developed by Ralph Keeney, Value-Focused Thinking (VFT). This methodology is then expanded to allow decision makers to specify not just constant weights to demonstrate their preferences, but an entire distribution. These distributions are then incorporated with the value of the attributes and the whole is simulated using Monte Carlo Simulation provided by Crystal Ball.

The result of incorporating these weight distributions into the model, is an empirical distribution for the value of an alternative. The alternative distributions can be compared in a number of ways to provide insight to the decision maker.

VALUE FOCUSED THINKING IN THE PRESENCE OF WEIGHT AMBIGUITY: A SOLUTION TECHNIQUE USING MONTE CARLO SIMULATION

1. Introduction

Countless decisions are made every day by people all across the world. Some of these decisions are made as a rapid response to immediate events allowing little or no time for analysis. Others are researched at length and made only after exhaustive study and debate. Whether made in the spur of the moment or after years of analysis, these decisions encompass a myriad of outcomes and consequences.

Throughout history, people have been making decisions. For almost every complex decision made, someone else is trying to help, guide or influence the decision and the decision maker. This aid comes from all areas of life. Religion, ethics, morality, science, politics, force and philosophy are only a few of the schools of thought that have attempted to guide the world in its decision making. While these sources may all provide good guidelines, their general rules and guiding principles do not provide insight into every decision situation. The growing complexity of our world has made it necessary to address our decision making from a more precise, organized and analytical perspective (Kaufman, 1968: 13).

In any decision, the Decision Maker (DM) must choose between at least two alternatives. For most of human history, the decisions facing mankind have been made

largely through intuition (Kaufman, 1968: 12). Decision Analysis (DA) seeks to aid the DM through structuring the decision problem clearly and choosing among competing alternatives (Howard, 1968b: 581-582).

1.1 Background

Even among the more analytical approaches to decision making, there are many possible techniques. Social and cognitive psychology offer an entire body of knowledge to the science of human decision making. Psychological research into social and cognitive processes may involve analyzing the specific mechanics used by the human brain in making a decision or what aspects of a person's life is most influential in shaping their decisions. Economics also provides techniques for making decisions involving fiscal or monetary situations. Operations research (OR) is another science involved in the pursuit of better decision making (Howard and Matheson, 1968: 21-26). Compared to psychology and economics, OR is a relatively new discipline which brings its own suite of tools and techniques to address decision making. Mathematical programming has been used to aid decision making involving optimizing some value in the face of constraints. Game theory addresses decisions made against direct opposition from some other person(s) or situation(s). Decision analysis seeks to provide insight into a decision based on the information available about the situation, the potential alternatives and the preferences and attitudes of the decision maker (Bunn, 1984: 1-8; Clemen, 1999: 5-8). Decision Analysis provides the larger canvas for the research in this thesis.

Decision analysis is concerned with identifying the logical and reasonable conclusions based on the problem structure, the alternatives, and decision maker preferences (Howard, 1983: 7). Decision analysis seeks to assist in making some impending decision. Through interaction with the appropriate decision makers, analysts and subject matter experts, DA attempts to accurately structure that decision and identify the relevant pieces.

Decision analysis has been applied to a number of decision problems in public policy, corporate decision making and personal choices (Keeney and McDaniels, 1999: 651; Keeney, 1992: 342, 372). The methods used to analyze decision situations in all areas come in many forms and focus on different aspects of the situation. The method employed and the factors under consideration are a direct result of the decision situation itself. Some methods address decision making with a single objective (maximizing expected profit for example) (Clemen, 1996: 19-21), while others concentrate on structuring problems where many diverse and often conflicting objectives are present (Keeney and McDaniels, 1999: 655). Decision analysis models can address problems with sequential decisions or when only a single decision needs to be made. These problems can include deterministic frameworks where all relevant parameters are assumed to be known (Keeney and McDaniels, 1999: 656-8) or can incorporate a number of uncertain events and variability in the outcomes. Finally, the goal of decision analysis is to provide insight into the decision being made and to provide some measure of or insight into the ranking of the competing alternatives.

Value-Focused Thinking (VFT) is one methodology designed to achieve the goal of providing insight to the decision maker. VFT was formalized by Ralph Keeney in his

1992 book *Value-Focused Thinking: A Path to Creative Decisionmaking* (Keeney, 1992) with the idea that values should be emphasized over alternatives. This approach seeks to identify the inherent values of the decision maker and the underlying structure of the decision to choose between competing alternatives or to even help identify new ones. Quantifying the relative importance of the various and often competing objectives in a VFT problem may be non-trivial. This relative importance is modeled in VFT as weights on the various attributes. The essence of these weights is in “value tradeoffs” (Keeney and Raiffa, 1976: 66). The decision maker must decide how much of attribute A he or she is willing to give up to get a certain increase in attribute B. These weights come directly from the decision maker and can be elicited in a number of ways. The methods for developing weights can impact their final values (Bottomley and Doyle, 2001 and Poyhonen and Hamalainen, 2001). Just as the elicitation varies, so do the exact interpretation of these weights (Choo, Schoner and Wedley, 1999). Despite the varied approaches to weight elicitation and the many interpretations of the weights themselves, none of these methods are designed to account for ambiguity. They all conform to the fundamental assumption that the weights can be determined absolutely and do not change.

VFT and many other decision analysis methodologies do, however, provide for the possibility that there may be some uncertainty in the performance of the alternatives on the measurement criteria. One source of uncertainty in alternative performance is a lack of information. The DM may not have all of the needed details to evaluate the performance of an alternative. A second source of uncertainty comes from evaluating a previously untried alternative whose performance can only be estimated ahead of time.

Regardless of the source, the presence of this ambiguity is a widely accepted circumstance (Clemen, 1996; Beauregard, 2001). Chapter 2 discusses how this ambiguity can be incorporated into a decision problem structure and eventually used to provide insight for the decision maker. One possible method for addressing this ambiguity is through the use of Monte Carlo simulation (Clemen, 1996: 413).

In 1964, J. M. Hammersley and D. C. Handscomb wrote *Monte Carlo Methods* in which they stated, “Monte Carlo methods comprise that branch of experimental mathematics which is concerned with experiments on random numbers,” (Hammersly and Handscomb, 1964: 2). This is often useful when only a small amount of real data can be obtained or the analytical solution is extremely difficult to derive or is intractable. In the DA context, if the ambiguity in some of the model parameters can be described through some set of random numbers or an underlying distribution, Monte Carlo methods can be used to develop probability distributions of the final value for each alternative (Clemen, 1996; Chacko, 1991).

1.2 Problem Statement

Decision analysis models, in their basic philosophy and structure, often assume a single DM (Buchanan, 2001; Howard, 1980: 198). Another assumption is that the weights or preference structure of the decision maker can be determined absolutely and without variation (Kassouf, 1965: 8; Lavelle *et al.*, 1997: 769; Shepetukha *et al.*, 2001:229). As mentioned in the previous section, VFT assumes that the weights used in the model are known and unchanging. In many situations, this assumption may not be

valid. The potential for a violation of this assumption is well known giving rise to the idea of sensitivity analysis (Winston, 1994). That is, analyzing the model to see if it is robust to changes in a given parameter. This is often a *post hoc* process accomplished after the primary analysis and not affecting the underlying structure of the decision (Felli and Hazen, 1999: 79). However, if the preference weighting is known to be ambiguous, it may be more accurately represented as a distribution and incorporated into the basic model as such. Chapter 2 explores, in more detail, when this situation might occur.

Furthermore, a disconnect has developed between the underlying theory of VFT and its actual application. VFT is often used to address problems faced by groups (Keeney and McDaniels, 1999). Unless a consensus in preferences is reached with certainty, some variability underlies the weighting. Even when dealing with a single decision maker, ambiguity concerning the DM's preference structure may exist (Moshkovich *et al.*, 2002). A DM's preferences may change over time, inducing variability into a decision structure meant to be robust against the timing of the decision (Fishburn, 1964: 20-21). An ambiguous preference structure may also be created through the existence of uncertainty or variability. Further, a decision maker may not have a complete understanding of his or her own preference (Fishburn 1964: 20-21, 84). This ambiguity on the part of the DM may adversely affect the outcome of the model if certainty is incorrectly assumed. A methodology designed to incorporate this ambiguity, regardless of the source, will allow greater flexibility in the possible structure of decision problems.

1.3 Research Direction

The research in this thesis demonstrates a process for incorporating weight ambiguity into VFT models. Through the use of Monte Carlo simulation, ambiguity can be modeled and the resulting decision problem analyzed empirically. A review of the current literature and practices provides the framework for the issues involved and establishes the premise for the need for this new technique.

1.3.1 Objectives.

Having identified a potential need for including weight ambiguity into the structure of a DA problem, this research has two objectives: develop a solution methodology and provide concrete examples of this methodology. First, this research develops and describes a methodology for incorporating weight ambiguity into the VFT framework. This methodology includes the use of Monte Carlo simulation to construct an empirical distribution of values for each alternative. These distributions can then be evaluated to determine the preferred alternative. Second, this methodology will be demonstrated through the use of two concrete examples. The first example is from a 1997 study attempting to develop an analytical solution to a problem of weight ambiguity (Lavelle *et al.*, 1997). The second example will come from the Information Assurance decision problem found in Lt Beauregard's thesis (Beauregard, 2001).

1.3.2 Approach

This study relies heavily on a literature search of decision analysis and decision making under ambiguity. Synthesizing this literature provides a map through the last half century of the discipline. During this time, the importance, use and elicitation of weights has gone through many phases. At times they have been thought of as merely “scaling constants” used to try to equate fundamentally different aspects of a decision problem; for example, how does a DM equate an increase of one mile per gallon of gas mileage to increasing a vehicle color from his or her second favorite to most favorite (Keeney, 1992). At other times, weights have very specifically been labeled as the decision maker’s specific relative preference among the competing objectives (Beroggi 1999).

Once the literature search has firmly established both a need and a possible avenue to fulfill this need, the methodology being proposed is described. The methodology presented in this thesis mirrors VFT up to the elicitation of weights. At this point, the discussion departs into Monte Carlo simulation. This approach also necessitates a new method for evaluating which alternative is “best.” While a number of possible techniques for comparing alternatives are presented, it is not the focus of this research to determine which of these, or some other technique not described, is of most use to the decision maker. The previously mentioned examples (Lavelle and Beauregard) have been chosen to demonstrate two things. The first example (described in detail in Chapter 2) compares analytical and simulation results for a given problem structure to establish the methodology presented as a valid and applicable approach to dealing with

the existence of weight ambiguity in a VFT framework. The second example (also described in Chapter 2) demonstrates the applicability and use of the methodology on a more complex decision hierarchy.

1.3.3 Scope and Assumptions

This research is not without its own assumptions and limitations. The literature search encompasses a lengthy discussion on the interpretation and elicitation of weights. The research does not, however advocate any given interpretation or go into any discussion of their relative merits. For more in depth comparisons and discussions of elicitation techniques, the 1999 study by Choo, Schoner and Wedley provides a survey of techniques that have been used (Choo *et al.*, 1999). It is not the intention of the author, or the point of this research, to enter into the debate on preferred interpretations or elicitation methods.

Similarly, this thesis does not discuss or evaluate the methods for eliciting subjective probabilities from a decision maker. Other researches have explored this area and developed a number of techniques and approaches (Benson *et al.*, 1995; Kahn and Meyer, 1991). This research assumes that these approaches are appropriate and sufficient to address the issues of weight ambiguity as used here.

Finally, this thesis is concerned with the VFT hierarchical structure and adapting that structure to accommodate a new dimension of ambiguity. This method uses the weighted, additive value model and its underlying assumptions. It does not include any process or guidelines for any other aggregation rules. While the basic approach

developed here may be appropriate for other multi-attribute decision analysis methodologies or systems, the author makes no claims to applicability outside of VFT. Expanding this approach beyond VFT, is an issue for future research.

1.4 Research Contribution

The techniques presented in this thesis represent a basic adaptation of a familiar method for analyzing multi-attribute decision problems. Through an expansion of weight elicitation methods already in practice, weight distributions can be assessed from the decision maker to correspond with the problem attributes. The presence of these distributions increases the complexity of the commonly used multi-attribute value function presented later in Chapter 2. While both a framework and an analytical solution technique (Lavelle *et al.*, 1997: 769) have been previously used to resolve very simple decision structures, it has limited application to more complex decision structures. Simulation can overcome these limitations in the analytical approach.

The key contributions of this research come from its expanded framework that allows for ambiguity in the weights and its development of a Monte Carlo Simulation approach to appropriately evaluate and analyze the value of alternatives in the presence of this ambiguity

1.5 Thesis Overview

This thesis explores the case in which the weights of a multi-criteria decision problem cannot absolutely be determined. Because of this ambiguity, at least some of

these weights must be specified by a range or distribution. The research uses Value-Focused Thinking to structure the decision, and employs Monte Carlo simulation to develop empirical distributions for the value of the alternatives.

Chapter 1 presented the background and framework to explain the question at hand. It also answers the research specific questions of objectives and assumptions. Chapter 2 expands this background into a literature search. Chapter 2 addresses utility theory and value preferences, the use of weights, properties of decision makers, ambiguity, alternative selection, Monte Carlo simulation and an overview of two specific decision examples that will be used later in the thesis. The methodology presented in Chapter 3 creates a simulation model for incorporating weight ambiguity in a VFT problem. The exact methodology used, and the resulting simulation, depends on which of the three separate weighting techniques presented is employed. Examples are presented in Chapter 4. Chapter 5 ends the thesis with summaries of the issues, the proposed technique and the examples. It also discusses some of the limitations of the research and ends with the final conclusions.

2. Literature Review

This chapter provides a review of the literature that is key to establishing the framework that is the basis for this research. This framework begins with a look at the broad issues and problems addressed by decision analysis. This broad overview is then narrowed to discuss utility theory, value theory and Value-Focused Thinking (VFT) specifically. After reviewing these theories and techniques, the chapter specifically addresses preferences and value tradeoffs a decision maker must identify. Decision making in the presence of ambiguity is addressed when these preferences and other model parameters are not exactly known. From this point, the focus shifts to Monte Carlo techniques and the use of simulation in solving problems that involve probability or uncertainty. The chapter ends with a discussion of possible formats used to provide decision makers with insight into the alternatives when ambiguity exists.

Throughout this chapter, several terms are used repeatedly. Some of these terms are used interchangeably, others are often thought to be interchangeable, but in truth have different meanings. The first set of definitions to look at involve: attribute, objective and criterion. These words are often used interchangeably in Decision Analysis literature (Canada and Sullivan, 1989: 211). However, while similar, their individual meanings are distinctly different. An *attribute* is some aspect of a decision that is important to a decision maker. If purchasing a vehicle, cost may be an important characteristic to consider. Cost is an attribute in the decision to purchase a vehicle. *Objective* represents direction of improvement or preference of the attributes (Canada and Sullivan, 1989:

211). In the vehicle purchase example, a low cost is the objective. A *criterion* is a standard or rule that guides decision making (Canada and Sullivan, 1989: 211). A criterion relates how much more attractive an option gets as the “level” of an attribute moves in the direction of the objective. How much better is a car with a cost of \$10,000 compared to one with a cost of \$15,000? This thesis does not use *criterion* or *objective* interchangeably with *attribute*, but rather according to their definitions just given. Other terms with possibly multiple interpretations are also used throughout this thesis. *Value* and *utility* have been mentioned and are addressed specifically in this chapter. As a central theme of this research, *ambiguity*, *uncertainty* and *variability* are presented and discussed. These final terms are specifically defined as they appear throughout this chapter.

To better explain many of the ideas in this chapter, a simple example can often illustrate a given concept. The notional vehicle example seen in the preceding paragraph is used in these cases. To provide consistency and context, the decision problem is the purchase of a new vehicle, with important attributes being *cost*, *functionality* and *aesthetic value*. The competing objectives are *low cost*, *high functionality* and *high aesthetic value*.

2.1 Decision Analysis

In 1968, Ron Howard stated, “Decision analysis has emerged from theory to practice to form a discipline for balancing the many factors that bear upon a decision,” (1968b: 1). Now, over 35 years later, the practice and application of decision analysis

has grown tremendously (Raiffa, 2002). Decision analysis comprises more than just theory and application, but an entire philosophy and world view (Howard, 1968b: 1; Howard, 1980; Raiffa, 2002).

The decision makers faced with these problems must make value judgments, for without value judgments, there can be no decision (Howard, 1968a; Keeney, 2002). Value judgments are necessary to identify appropriate value trade-offs (Keeney, 2002: 936). In this respect, a good value trade-off is one that accurately represents a decision maker's values concerning a decision. These value trade-offs, and the underlying value judgments, drive the decision process as the true basis for comparing alternatives (Keeney and Raiffa, 1976: 66-69). Alternatives are ultimately the cause for a decision in the first place. Without at least two alternatives, the course of action is simple, and no real decision is required; one simply follows the course of action available. Competing alternatives most often involve more than one attribute. The decision maker's objectives for these attributes are often in conflict (Watson and Buede, 1987: 19). That is, to increase the desirability of an alternative in one attribute will often require lowering the desirability in another. This situation is most clearly seen when discussing the cost of an item. Generally, higher quality and functionality and lower cost are competing objectives. It is the competing nature of objectives that often makes it difficult to choose among alternatives.

Decision analysis separates a decision into its component pieces to better provide for a rational decision. A rational decision is one that is logically consistent. A rational decision maker is one whose decisions flow logically from a set of given values and preferences (Howard, 1980: 181-182; Kassouf, 1970: 1-8). Decision analysis is, by

nature, a normative process, that is it seeks to dictate how decisions **should** be made rather than describing how they are being made (Howard, 1980: 181; Kassouf, 1970: 1-4). One question then becomes, why do we need to tell a decision maker what should be done if he or she makes rational decisions? Unfortunately, most decision makers are not rational when presented with new problems (Howard, 1980: 181). It is the ability of decision analysis to assume rationality and then express how a decision should be made (to maintain that rationality) which is its true “power” (Howard, 1980: 181). By following the actions prescribed through decision analysis, it is possible to improve on our natural decision making abilities and make better decisions (Howard, 1980: 181-186; Kassouf, 1970: 1-4).

2.1.1 Utility

One very common normative approach to decision analysis is utility theory.

Watson and Buede provide a basic definition of utility theory as:

The concept of a numerical measure to describe the value of alternative choices has come to be referred to as *utility theory*, with the *utility function* being the numerical measure itself. (1987: 21)

“Alternative choices” broadly refers to any possible outcomes or consequences resulting from some decision. For any two possible outcomes (x_1 and x_2), a number [$u(x_1)$ and $u(x_2)$] can be assigned to each outcome such that x_1 is preferred over x_2 if and only if $u(x_1)$ is greater $u(x_2)$:

$$x_1 \succ x_2 \Leftrightarrow u(x_1) > u(x_2)$$

where \succ means “is preferred to” and the numbers assigned [$u(x_1)$ and $u(x_2)$] represent the utility of x_1 and x_2 respectively (Fishburn, 1970: 9).

What the definition at the beginning of the preceding paragraph does not explain is that utility theory is concerned with the uncertainty in a decision problem. More specifically, it addresses the uncertainty associated with the possible outcomes of a decision. In the impending vehicle purchase, the exact cost of any given alternative may not be known exactly. The cost may fluctuate based on the day of purchase, the sales person available or the negotiating skills of the decision maker. The outcome is uncertain. It is very important at this point to emphasize this use of uncertainty. Utility theory addresses the decision maker’s reaction to uncertainty in outcomes, not preferences (Kahn and Meyer, 1991). The decision maker’s reactions, or attitudes, towards uncertainty in outcomes are characterized by his or her risk preference, that is, whether he or she is risk seeking, risk averse or risk neutral (Watson and Buede, 1987: 21; Kassouf, 1970: 36).

The decision maker’s risk preference can be described through the use of a utility function (Watson and Buede, 1987: 21). The relationship between preference and the utility function is similar to the utility/preference relationship described at the beginning of this section. There exists a real valued function, $u(\square)$, such that for every possible outcome, x , in X :

$$x_1 \succ x_2 \Rightarrow u(x_1) > u(x_2). \text{ (Fishburn, 1970: 9)}$$

The real valued function, $u(\square)$, is the utility function for the set, X , of possible outcomes. To assess the utility function, the decision maker is asked a series of questions about his

or her preference regarding the uncertain outcomes. The decision maker states his or her preference between a certain outcome and the uncertain outcome of a lottery. These lotteries are often in terms of monetary value (Keeney and Raiffa, 1976: 143). The decision maker can choose a certain amount of money $\$y$ or can choose the lottery with a probability, p , of receiving $\$s$ and a probability, $(1 - p)$, of receiving $\$w$. When the decision maker is indifferent between the certain outcome and the lottery, $\$y$ represents the *certainty equivalent* (Fishburn, 1970: 117; Howard, 1968b: 584). The utility function can be identified through assessing the certainty equivalent at several points within the set of possible outcomes (Keeney and Raiffa, 1976: 204-206). Any of the parameters ($\$y$, $\$s$, $\$w$ and p) involved in the determination of a certainty equivalent can be varied to identify the point of indifference. Which parameters to vary, the exact structure of the lottery and the way in which the questions are posed to the decision maker are all dependent on the decision maker, analyst, and the current decision situation (Bunn, 1984: 32-33; Clemen, 1996: 469-480; Fishburn, 1970: 117-119; Keeney and Raiffa, 1976: 142-208). The expected value of this utility function is known as the expected utility of the possible outcomes.

In 1738, Daniel Bernoulli wrote, "...the *value* of an item must not be based on its *price*, but rather on the *utility* it yields. ...the utility...is dependent on the particular circumstances of the person making the estimate," (Bernoulli, 1954: 24). He was writing about how to reconcile the apparent inconsistency in how people value money (Watson and Buede, 1987: 19-21). Utility has since been used to help describe economic behavior and ethics. Keeney and Raiffa declared that von Neumann and Morgenstern created the axioms and foundations for the utility theory used widely in decision analysis (Keeney

and Raiffa, 1976: 131, 283). Since then, others have built on their work, with significant accomplishment in the field by Peter Fishburn (Fishburn, 1970) and later by Ralph Keeney and Howard Raiffa (Keeney and Raiffa, 1976) (Watson and Buede, 1987: 21). Although this thesis is chiefly centered around work on value models, a general understanding of utility theory will help in understanding the discussion on ambiguity later in this chapter.

2.1.2 Value

In 1987, Watson and Buede made a very clear distinction between value and risk preferences (21). The von Neumann-Morgenstern utility functions discussed in the previous section are used in cases of uncertainty and represent the decision maker's risk preferences (Watson and Buede, 1987: 21). *Value* is a measure of worth to a decision maker. Value judgments, preference structures and value tradeoffs from the decision maker all contribute to *value* as a measure of worth (Keeney and Raiffa, 1976: 66). Value measures the worth of an outcome to a decision maker regardless of the probability that the outcome will be realized. That is, value is not concerned with uncertainty, but simply the preferences of the decision maker (Watson and Buede, 1987: 21).

In *Decisions with Multiple Objectives* Ralph Keeney and Howard Raiffa make a “digression” to draw a clear difference between utility functions as created by von Neumann and Morgenstern and the utility function often used by an economist. Specifically, Keeney and Raiffa discuss decreasing marginal utility and explain that because this has no probabilistic aspects any expected utility calculated from the

decreasing marginal utility curves is useless because “expected utiles” has no real meaning (Keeney and Raiffa, 1976: 150). They further describe exactly what is meant by a decreasing marginal utility function:

As an example of our economist’s *utility function* with decreasing marginal utility, suppose we considered 8 utiles as the utility of one day of skiing, 14 utiles for two days, 18 utiles for three days, and so on. Then we could say the first day is worth 8 utiles, the second an additional 6, and the third another 4. The marginal utility of each additional day of skiing is decreasing. However, if we had a choice between two days of skiing for sure or a lottery yielding either one or three days with equal likelihood, we could not say which option should be preferred using the *utility function*. This is so even though the expected number of utiles is 13, whereas it is 14 for the sure two days skiing. The concept of “expected utiles” has *no* meaning. (Keeney and Raiffa, 1976: 150)

In this example, Keeney and Raiffa explain how the absence of uncertainty in the decreasing marginal utility function prevents it from having any true meaning as a von Neumann-Morgenstern utility function. This example, for the same reason, is excellent in describing the difference between a value function and a utility function. Further, the meaning and use of the decreasing marginal utility is a fair analog to the meaning and use of value. In the example, utiles are measured in whole units and describe worth to the decision maker. If these units were normalized to sum to 1 (assuming that three days skiing is the maximum possible outcome for this situation), then one day of skiing would now have a utility of 0.44 ($8/18$), two days of skiing would have a utility of 0.78 ($14/18$), and three days would have a utility of 1 ($18/18$). However, in the context of decision analysis as described so far, these are not utilities because, as Keeney and Raiffa explained in the original example, they do not incorporate any uncertainty. Instead, the decreasing marginal utility has captured the decision maker’s relative value tradeoffs among the possible outcomes, i.e. the number of days skiing. The *utility* represents the

value of each outcome and the decreasing marginal utility function can be thought of as the value function. In this case, any alternative that provides all three days of skiing provides the maximum value (1) for the objective “skiing”. If, however, an alternative only provides 2 days of skiing, the value for the objective “skiing” for that alternative would only be 0.78. From this, it is clear that the decision maker’s “utility,” or worth, in the absence of risk or uncertainty in the outcome is his or her value. Throughout this thesis, *utility* is used to represent the decision maker’s preferences in the presence of uncertainty in the outcomes, i.e. risk preference, and *value* is used to represent the decision maker’s preferences when there is no uncertainty in the outcomes of an alternative.

The research in this thesis is based on the weighted, additive value model in Equation 1:

$$V_j = \sum_{i=1}^k w_i v_i(x_{ij}) \quad \forall j = 1 \dots p \quad (1)$$

where:

- $V_j \equiv$ the overall value of alternative j
- $w_i \equiv$ the weight of measure i
- $v_i \equiv$ the single dimension value function for measure i
- $x_{ij} \equiv$ the score of alternative j on measure i
- $k \equiv$ the total number of measures
- $p \equiv$ the total number of alternatives

In this equation, the value of an alternative is the weighted sum of the values obtained from the measures created for each alternative. The processes for developing this equation are presented in section 2.2.

2.2 Value-Focused Thinking (VFT)

Value-Focused Thinking (VFT) was formalized by Ralph Keeney and presented in his book *Value-Focused Thinking, A Path to Creative Decision Making* (Keeney, 1992). VFT is a hierarchically structured, multi-attribute decision analysis methodology that emphasizes value over alternatives. Many decision methods focus quite heavily on the alternative solutions available for a given decision (Keeney and Raiffa, 1976; Olson *et al.*, 1999). This does not seem abnormal when one considers that, “decision making usually focuses on the choice among alternatives,” (Keeney, 1992: 3). Instead, VFT structures the decision problem based on the decision maker’s values. Ralph Keeney wrote, “Values are what we care about,” and then, “Values are more fundamental to a decision problem than are alternatives,” (1992: 3). In this respect, VFT is more than a specific methodology. It is a philosophy as well as a general approach to the science of decision making and the process of providing insight to the decision maker.

“Values of decisionmakers are made explicit with objectives,” (Keeney, 1992: 33). A *fundamental objective* provides the decision context, the value involved and the direction of preference (Keeney, 1992: 34). Values can be further refined to means objectives. *Means objectives* represent the specific *means* by which the fundamental objectives can be achieved (Keeney, 1992: 34-35). The degree to which these means objectives are achieved is determined using a measure or *measurement scale*. Please note that while a measurement scale is used to define an *attribute* (Keeney and Raiffa, 1976: 32; Kirkwood, 1997: 12) it is also used in a broader sense in the general literature (and as noted at the beginning of this chapter) to indicate all of the “factors of importance”

involved in a decision problem (Lavelle *et al.*, 1997: 769). In the VFT context, the “factors of importance” are the decision maker’s values. To continue with this broader usage, this thesis refers to values, fundamental objectives and means objectives collectively and alternatively as attributes.

As a multi-attribute technique, VFT breaks the decision problem into its component pieces based on the values of the decision maker. The vehicle example already introduced in this chapter is an example of a multi-attribute problem. The decision context is the purchase of a new vehicle. The DM’s values in the purchase of a new vehicle guides the VFT process. These values are structured into a hierarchical form (Kim and Han, 2000: 79). This form is also referred to as a hierarchical attribute tree, a hierarchical value tree or a value hierarchy (Keeney, 1992; Kim and Han, 2000: 79; Poyhonen, 1998: 7). Figure 1 provides an example of a value hierarchy:

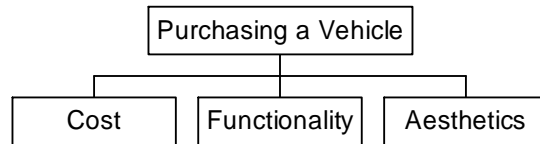


Figure 1: Notional Value Hierarchy for a Vehicle Purchase

In Figure 1, “Purchasing a Vehicle,” is the decision of interest. For this decision, the DM values *cost*, *functionality* and *aesthetics*.

VFT, as originally presented by Ralph Keeney, is designed to deal with uncertainty and is rooted in utility theory (Keeney, 1992: 129-141). The mathematical underpinnings of VFT are in development and use of utility functions (Keeney, 1992: 129-154). Keeney refers to the resultant models as value models (Keeney, 1992: 129-132). Why use *utility* functions to develop a *value* model? Keeney provides three

reasons. First, the use of utility functions allows him to address the presence of uncertainty and risk. Measurable value functions do not address risk attitudes. Second, concentrating on utility functions allows for a more concise discussion. Finally, “the concepts and procedures for utility functions and measurable value functions are analogous,” (Keeney, 1992: 132) Why a value model instead of a utility model? Utility functions are created using the decision maker’s value trade-offs (which, in turn, came from the DM’s value judgments) (Keeney and Raiffa, 1976: 220-222). As mentioned earlier in this section, VFT is more than just a methodology, it is a philosophy for approaching decisions. In cases with no uncertainty, the value principles discussed in section 2.1.2 apply.

The general VFT methodology and philosophy can be applied to value models as well as utility models. In cases where the value adapted VFT methodology is used, the general process follows along the following lines. 1) Identify and structure the objectives and criteria important to a decision maker for the decision situation at hand (Keeney, 1992: 55-98). This structure often takes the form of an objectives hierarchy that continues to break down the component parts of an objective until a final end measure can be found. 2) Develop measurement scales for each of the final end measures identified in step 1 (Keeney, 1992: 3) These measures, which are often composed of scales with different units and vastly different ranges, are assigned a value according to some function: $v_i = f(x_i)$, $\forall i = 1 \dots k$ where x_i is the resulting score of some alternative on the measure associated with a single attribute i , and $f(x_i)$ is a monotonically increasing or decreasing function whose range is between zero and one (Kirkwood,

1997). These functions, known as Single Dimension Value Functions (SDVF), are created for all k measures. This mapping translates scores from measures into value. The value is measured from 0 to 1 with higher value being more preferred. 4) Weight the competing objectives based on the preferences of the decision maker. 5) Score each of the identified alternatives on the measurement scales and convert them to values. 6) Aggregate the weighted values from each of the lowest level attributes and determine final values for the alternatives. 7) Rank these alternatives and provide insight to the decision maker.

A decision analysts greatest value to the decision maker is in the ability to properly structure a decision situation (Howard, 1986b: 2). By focusing on values rather than alternatives (Keeney, 1992: 3) the analyst is able to accomplish this regardless of the current set of alternatives. A properly structured value model from the decision maker can be used not only to rank existing alternatives, but also identify shortfalls and strengths in these alternatives based on the objectives and preferences of the decision maker (Keeney, 1992: 9). If these alternatives are unsatisfactory as a whole, the structure can be used to develop potentially better alternatives. Emphasis on values and the given stepwise procedure has also been applied using value functions rather than the utility functions provided in the original VFT framework but is still generally referred to as Value-Focused Thinking (Beauregard, 2001).

2.3 Weights

This section focuses on that portion of VFT concerned with a decision maker's preferences and value trade-offs. These preferences and value trade-offs are represented as weights in a value model. If the weight for a given objective is higher than the weight of another, that objective is considered to be more important to the decision being made based on the decision makers values and preference. The potential interpretation and elicitation of these weights is quite diverse and has an impact on how the problem is structure and the insight provided (Choo *et al.*, 1999; Poyhonen, 1998). Some of these interpretations include the marginal contribution per unit of the objective in question, discriminating power of the objectives, voting values, relative functional importance and others (Choo *et al.*, 1999).

Regardless of their interpretation, value trade-offs must be tied to the range of the raw scores for the measures to which they are linked (Keeney, 2002: 940-941). This can be illustrated with the vehicle example. In the vehicle example, cost may be considered in general more important than aesthetics by the decision maker and one may expect to see cost weighted higher than aesthetics. However, if the range in the cost of each of the competing alternatives is quite small, the decision maker may feel that the difference in value of decreasing cost from its highest to its lowest may not be overall as important as the change in aesthetics. Of course, exactly what "quite small" entails will be up to the decision maker. In this case, cost will have a lower weight reflecting the small range among alternatives and smaller value the decision maker puts on a change of cost within that range. Misinterpretation of these weights, no matter how they are applied, can often

lead to incorrect rankings of alternatives and provide faulty insight into the problem (Choo *et al.*, 1999).

Not only do the weights have potentially differing meanings, but they can be elicited from the decision maker in a number of ways. These elicitation methods are inextricably linked to their interpretation. One thing that weight elicitation methods and interpretations have in common is that they result in constant weights (Choo *et al.*, 1999). This philosophy of constant weights stems from an early concept that the weights fall into the category of decision variables and as such can be set to any level by the decision maker (Spetzler *et al.*, 1972). As described in the Problem Statement in section 1.2, decision analysis methodologies assume that these weights can be exactly specified by the decision maker.

Among the various weight elicitation methods used, the exact information obtained from the decision maker changes. In some cases, the information provided by the decision maker is not the weights that will be used to solve Equation 1 but rather other preference structures that can then be used to calculate these weights (Poyhonen and Hamalainen, 2001: 569-572). Regardless of the method used, it is generally accepted, to provide consistency in application, that the weights (directly or calculated) must sum to 1 (Poyhonen and Hamalainen, 2001: 570; Kirkwood, 1997: 70). Whether weights are elicited directly or are derived from other preference statements, they can be obtained either locally or globally. A value hierarchy is grouped into tiers. Each tier represents further breakdown and delineation of the “parent” attributes in the tier above. Weights can then be elicited based on each tier (hierarchical weighting or local weights). Weights could also be elicited directly from the lowest level attributes in the hierarchy

(non-hierarchical weighting or global weights) (Poyhonen, 1998: 8). In the case of hierarchical weighting, the final weights used in the calculation of Equation 1 are found by multiplying the local weights to the end of each branch. In non-hierarchical weighting, the global weights are used directly in the calculation of Equation 1.

2.4 Direct Weighting

Direct weighting involves having the decision maker provide the exact numerical weight (from zero to one) for each of the attributes in question. This is often accomplished by telling the decision maker he has 100 points and must allocate all of them among the objectives. This is also sometimes referred to as the “100 balls” method. These direct weights can easily be seen as a direct representation of the value of the attribute relative to the whole (Poyhonen *et al.*, 2001: 571). In the vehicle example, a decision maker may say that cost is 40 points (40 out of a possible 100 is 40% of the decision) and therefore has a weight of 0.4. This weight has meaning because we know that the total sum of the weights must be 1. If the sum of the weights were something other than 1, 0.4 would no longer mean 40% of the decision (it would not account for 40 points out of 100). The number provided by the decision maker is the percentage of the decision encompassed in that attribute directly. Furthermore, the information provided about any given weight offers no direct indication of the weights of the other attributes. Only through their relationship to the whole (one) can relative proportionality be determined.

Direct weighting can be used locally or globally. When used locally, it assesses value trade-offs within a given branch and tier and assigns some portion of the whole value to each of them. When used globally, all final branch ends are compared for value tradeoffs and weights are assigned accordingly.

2.5 Independent Scale Weighting

Another method that has been used is independent scale weighting (Keeney and McDaniels, 1996; Lavelle *et al.* 1997). With the independent weighting technique, the information obtained from the decision maker relates the importance of each attribute to some independent scale that may or may not be directly related to each attribute. Whether directly related or not, the same scale is used for all attributes. Some scales that have been used are money (Keeney and McDaniels, 1996) and simply importance (Lavelle *et al.*, 1997: 772). In each case, an attribute is given some value along the scale used. To provide weights that sum to 1, the value of each attribute along the independent scale is divided by the sum of the values of all of the attributes. This normalization procedure relates the actual information obtained to the weights used.

In the car example, an independent scale from 1 to 10 might be used to determine the importance and subsequent weights of the three attributes. The decision maker is instructed to identify the importance of each attribute on this scale with 10 being very important to the decision maker and 1 being of very little importance. The decision maker may decide that cost is very important (assigns a 10), functionality somewhat important (assigns a 7) and aesthetics of little importance (assigns a 4). The sum of these

three values is 21. After normalizing, the actual weights would be $\text{cost} = 10/21 = 0.477$, $\text{functionality} = 1/3 = 0.333$, and $\text{aesthetics} = 4/21 = 0.190$.

Independent scale weighting does not provide a value that is directly relative to either the whole or among attributes. Knowing that cost gets a 10 in importance on a scale of 1 to 10 does not tell us how much of the value of the whole decision is captured in the cost until the values are normalized. There is no direct interpretation of “10” with respect to the whole decision. Similarly, there is no direct relationship among attributes by knowing only a single attribute value. Although cost is 10, there is no way to know how that might compare to other attributes unless their value is known. If the other attributes each get a 1, then cost is much more important relatively. If, however, the others are also given a 10, then cost, while important to the decision (as seen by receiving the highest score on the scale) is no more important than the others. But, in both cases, it is necessary to have more information than just the value of cost.

2.6 Swing Weighting

Swing weighting is the third method reviewed in this thesis (Kirkwood, 1997; Keeney, 1992). There are two variants to this method. The first variant is to take each measure independently and consider what increase in value of the alternative as a whole would result from swinging that measure from its lowest possible score to its highest. This is done with all attributes and the resulting increases in value are ordered from least to greatest. Each increment is then assessed as some multiple of the least important increment. These values are normalized to provide the weights for the objectives

(Kirkwood, 1997: 68-72). The second variant involves taking the objectives in pairs (*A* or *B*) and holding all other objectives constant, deciding which of the pair, *A* for example, the decision maker would rather swing from its lowest to its highest value. The decision maker must also provide the strength of this preference (Kirkwood, 1997: 68-72). This is done with all pair wise comparisons and the results normalized. If inconsistencies exist in the comparisons, they can be discussed with the decision maker and resolved. This method obtains information from the decision maker that relates one attribute to another. While the attribute of lowest importance is generally used as the baseline, this is not necessary.

2.7 Decision Making in the Presence of Ambiguity

“Most, if not all, decisions are made under uncertainty ...” (Wallace, 2000: 20). The decisions faced by individuals, groups and organizations encompass ambiguity as a major aspect (Watson and Buede, 1987: 11; Howard, 1983:7). As discussed earlier, uncertainty is considered in terms of the performance of an alternative or the possibility of outcomes. This uncertainty can stem from imprecision of measurement or the uncertainty inherent in trying to predict the future performance of a system not yet in use. This research, however is chiefly concerned with the situation in which the weights for the attributes are not known precisely.

At this point, it is important to define three terms: *uncertainty*, *variability*, and *ambiguity*. *Uncertainty* is a lack of knowledge about the true state of some quantity (Frey, 1993: 2). If the true state of some variable is known, but changing according to

some other variable (often time or individual samples), then that quantity exhibits *variability*. In this research, these two aspects of imprecision in information are collectively referred to as *ambiguity*. The term ambiguity will be used when it is not necessary to differentiate the nature of the imprecision, when both types are being referred to, or when this information may not be known. Uncertainty or variability are used in those cases in which their specific definitions apply and are relevant. Regardless of the source or cause, it is very beneficial to have a decision analysis model that can account for the presence of ambiguity.

2.7.1 Uncertainty in Measures

Uncertainty in the performance of an alternative or the possible outcomes has been a part of decision analysis models from the very beginning and almost all normative techniques are designed to account for it. Uncertainty in the performance of alternatives refers to the incomplete knowledge of how a given alternative will score against some measure or the natural uncertainty in the performance the alternative will exhibit through time (Wallace, 2000: 20). This uncertainty can also rise from the uncertainty of the environment in which the decision will ultimately be made (Clemen, 1996: 2).

2.7.2 Decision Makers

The nature and role of the decision maker can often affect the structure and technique applied. The most common use of “Decision Maker” or DM in the decision analysis framework is that of a single decision maker (Buchanan *et al.*, 2001; Lavalley,

1978). This idea of a single decision maker allows analysts to make simplifying assumptions in their models by not considering the interaction of the various parties to the decision and their possibly competing preference structures (Keeney and Raiffa, 1976: 516). However, it is becoming more apparent that many decisions involve not just a single decision maker, but a group that is responsible for making a decision. Another interpretation of the “decision maker” involves looking at stakeholder groups (Keeney and McDaniels, 1999). Even if a single person is responsible for making a decision, that person must often take into account not only their own preferences, but those of other stakeholders in the process.

Consensus is often the goal of group decision making (Ellis and Fisher, 1974: 141-143). Coming to an exact consensus through discussion may be a very lengthy and difficult process and may leave many members in the group feeling as though their views are not adequately represented (Ellis and Fisher, 1974: 235-246). Another issue involved when dealing with groups and group consensus is that a group will over time develop its own distinct identity (Rothwell, 1992: 183-187). This identity and the consensus it forms may not truly represent the preferences of the individual members.

2.7.3 Weight Ambiguity

Recent trends in decision analysis have pointed out the possible fallacy of assuming that the weights placed on objects are known with absolute certainty (Kelley and Thorne, 2001; Levary and Wan, 1998; Kim and Han, 2000; Lavelle *et al.*, 1997; Stewart, 1995). As a fairly recent relaxation of the assumption that preferences and value

trade-offs can be exactly determined, there is not yet a lengthy body of literature. There are several potential sources of ambiguity in objective weights. Even a single decision maker with plenty of time may not be able to exactly define and articulate his or her preferences. Researchers in this area have used terms like “partial information” and “incomplete knowledge” to describe this phenomenon, but regardless of the terminology, the concept is that the decision maker may not be able to completely resolve the complex changing environment in which the decision must be made (Kim and Han, 2000).

Far more likely than uncertainty within a single decision maker is the variability inherent in group processes and multiple stakeholders. The Keeney, McDaniels study for BC Gas, clearly illustrates the wide range of responses for value tradeoffs from the various stakeholders (Keeney and McDaniels, 1999). At times the variations involve differing orders of magnitude. It may be very difficult to represent these ranges as a single constant value. Possible solutions to this problem have included analytical solutions using greatly simplified MAVT models (Lavelle *et al.*, 1997) and a simulation approach based on building a probability of specific alternative rankings used by Kelly and Thorne (2001). Simulation was also used in conjunction with a descriptive decision analysis model, AHP, by Levary and Wan (1998). Kim and Han used a recursive mathematical programming approach to resolve uncertainty (2000).

2.8 Monte Carlo Simulation

Monte Carlo Simulation is a technique used when ambiguity exists in a system to develop an empirical probability distribution for some measure of the system (Kalos and

Whitlock, 1986: 129-130). More specifically, it is a numerical technique in which various input parameters are specified by some probability function. An output measure defined by some function of these random input variables will itself be a random variable. While probability theory can possibly provide an analytic solution to this function of random variables, it may be extremely difficult. In some cases there may not exist a closed form solution. To account for these difficulties, Monte Carlo techniques make a random draw from each of the input distributions underlying the model and use the result in the governing function for a given output variable. This process will provide one sample point from the distribution of the output variable. When a series of random draws is conducted many times in an established experimental design, an empirical distribution of the output variable can be constructed. This empirical distribution can then be used to address questions about the output variable.

Monte Carlo simulation is a logical alternative to standard expected utility theory (Clemen, 1996: 410-414). Instead of merely taking the expected utility of an alternative (in the case of risk preference and uncertainty in outcomes), analysts can use this technique to construct an empirical distribution of the overall utility of the alternative. Alternatives can then be ranked or investigated based on their distributions as well as their expected values. When ambiguity in the weights is present, Monte Carlo simulation can be used in the same manner it is used in the presence of ambiguity in the outcomes.

Monte Carlo techniques were first used to conduct numerical integration involved with the development of the atomic bomb (Hammersly and Handscomb, 1964: 6-9). In this and other deterministic applications, Monte Carlo simulation is used to estimate the value of integrals and other mathematical problems where no closed form analytical

solutions exist. In the 1950s, the growing development of operations research provided a new set of problems that could benefit from repeated sampling (Hammersly and Handscomb, 1964: 8-9).

Using weight distributions to replace point estimate weights in the weighted, additive value equation does not over complicate the equation itself. However, it's the solution to this equation that requires multiple integration techniques (Lavelle *et al.*, 1997: 773-779). This integration may, through more complex decision structures, become quite difficult and unwieldy for practical applications. This problem is highlighted further when using direct weighting in which the weight distributions are not independent due to their need to sum to 1. Monte Carlo simulation is useful in those cases where a problem can be formulated theoretically, but not solved that way (Hammersly and Handscomb, 1964: 3). Rather than develop the analytical solution from the value equation which now contains random variables, Monte Carlo simulation is used to sample from each of those random variables in turn and computes the solution to the equation. Repeated random draws create an empirical distribution for the final value of the alternative. This can be used when no analytical solution is possible.

The resulting empirical distribution need not approximate a normal distribution in order to derive inferences about certain distribution parameters (Davison and Hinkley, 1997: 25). Once an empirical distribution has been created for the final value of each alternative, they can be compared to determine the most attractive decision to the decision maker.

2.9 Selecting Alternatives in the Presence of Ambiguity

When Value-Focused Thinking is used to establish single constant values for alternatives in question, ranking the alternatives is a fairly straightforward matter.

Human judgment and sensitivity analysis can be used to provide insight into alternatives whose overall value are relatively close. These alternatives can be evaluated further, but the results will again be some point estimate of value that can be easily ordered.

Selecting the best alternative is not, however, as simple when the values of the alternatives are represented by some probability distribution. Chapter 3 will discuss the use of dominance and statistical tests to help determine which alternatives should be selected over others. Although they are included, the relative merits of the techniques used to compare alternatives is not addressed in this thesis.

2.10 Examples

As a new adaptation of VFT and an untested method for employing Monte Carlo Simulation, it is important to provide some level of justification that the methodology proposed in this research is appropriate. To this end, two examples are presented. The first example involves finding a location for a new airport. This first example was used to incorporate weight ambiguity and proposed an analytical solution techniques to arrive at distributions of value for each alternative. By simulating this example, a comparison can be made between the simulated results and the analytical results. This comparison will show that the simulated results closely approximate the analytical results. The second example involves selecting computer systems to maximize information assurance.

This example was chosen to demonstrate the methodology and its results on a more complex value hierarchy.

2.10.1 Mexico City Airport Citing

This first example illustrates the comparison of results between the analytical and simulated results. In 1997, the journal *Computers and Industrial Engineering* published an article by Lavelle *et al.*, “A Method for the Incorporation of Parametric Uncertainty in the Weighted Evaluation Multi-attribute Decision Analysis Model.” This article presents a multi-attribute utility model based on multi-attribute utility theory (MAUT) described by Keeney and Raiffa (Keeney and Raiffa, 1976: 436-472). The Lavelle *et al.* model is referred to as the weighted evaluation (WE). The WE model is a simplified adaptation of the weighted additive model shown in Equation 1. The simplifying assumptions are: a) attribute independence, b) linear utility functions and c) additivity of multiple attributes (Lavelle *et al.*, 1997). The assumptions of attribute independence and additivity of multiple attributes allows the authors to use the additive utility function:

$$Z_j = \sum_{i=1}^k W_i R_{ij}, \forall j = 1, \dots, p \quad (2)$$

where Z_j = the weighted evaluation of alternative j , W_i = the weight of attribute i , R_{ij} = the rating of alternative j on attribute i , k = the total number of attributes, and p = the total number of alternatives (Lavelle *et al.*, 1997). Each of the alternatives is rated on an independent scale of 0 to 100.

This WE model is then converted to a probabilistic weighted evaluation (PWE) by replacing the constants used in the Equation 2 with random variables. Lavelle *et al.*'s PWE method uses uniform, triangular and beta distributions to model uncertainty in both the weights on the attributes and the ratings for each alternative on each attribute. These distributions are used in calculating the first three central moments of the resulting distributions for the alternatives. The first central moment, the mean, is analogous to the point estimate results from the WE model. The second central moment, variance, looks at the spread of the distribution around this mean. Finally, the third central moment, skewness, gives a measure of how symmetrical the resulting distributions are. These three moments are then used to develop a normal approximation for the distribution of each alternative. The PWE does not ensure that the weights sum to 1.

2.10.2 Information Assurance

The second example used later in chapter 4 is the Information Assurance problem developed by Lt Joe Beauregard (Beauregard, 2001). Lt Beauregard modeled the effects on information assurance protection for a given computer security system using the hierarchy in Figure 2.

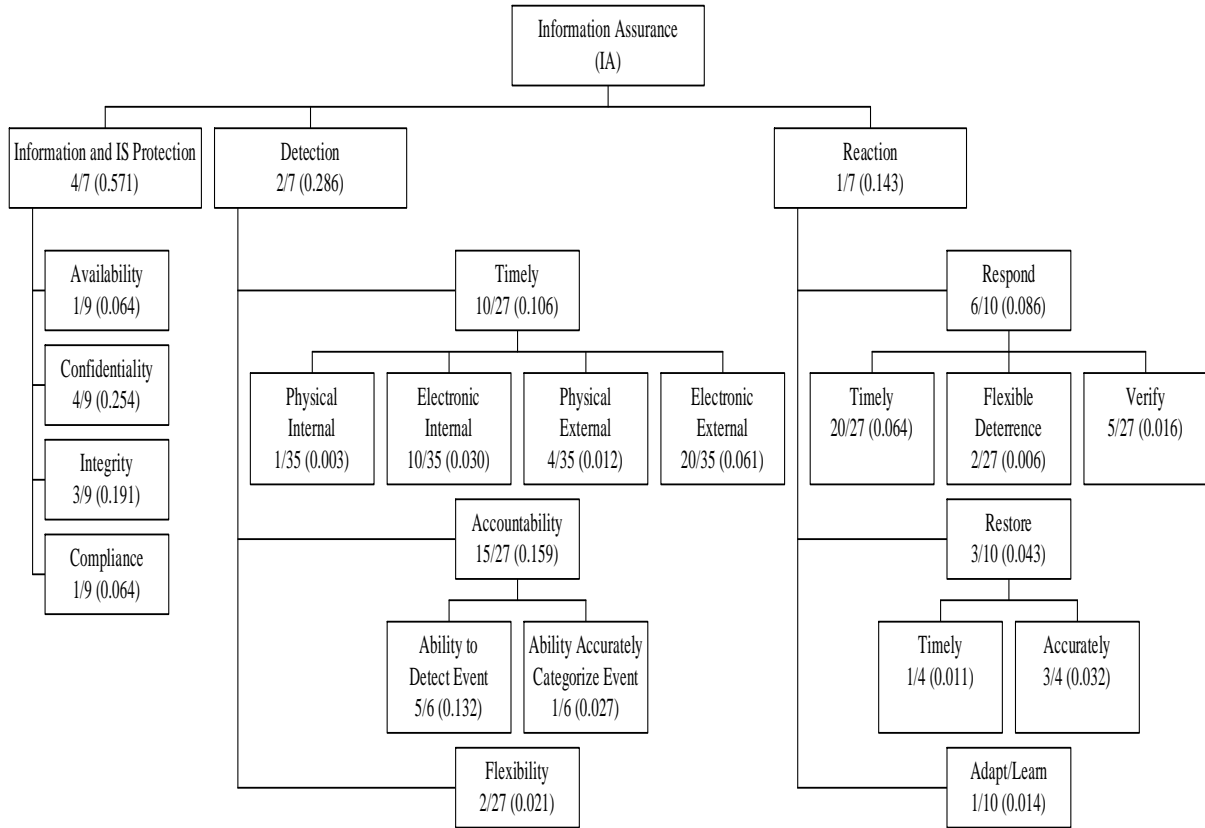


Figure 2: Weighted Information Assurance Hierarchy (Beauregard, 2001)

Figure 2 provides the name of each objective and sub-objective. It also includes two weights. The first weight given is the local weight. These local weights are given as a fraction based on the swing weight coefficients. The numerator of this fraction is the swing weight coefficient for the objective or sub-objective in question. The denominator is the sum of the swing weight coefficients in each branch and tier. The swing weight coefficients shown (the numerators) are assumed to be uncertain later in this thesis. The second weight, shown in parentheses, is the calculated global weight. This is the product of the local weight of the objective or sub-objective in question with the local weight of each successive parent objective or sub-objective above it.

This decision problem also included more than one measure for some of the final sub-objectives. In this research, the weights on those final measures are modeled as the point estimates given. This research only models ambiguity on the swing weight coefficients used in the local weights of the objectives and sub-objectives.

This chapter presented the framework for value models, VFT and the existence of weight ambiguity. It also discussed some of the implications of weight ambiguity and some possible solution methods. Two examples that will be used in Chapters 3 and 4 were also presented. The proposed methodology using Monte Carlo simulation to incorporate weight ambiguity in a VFT decision problem is further explained in Chapter 3.

3. Methodology

The methodology presented in this chapter diverges from Value-Focused Thinking in the determination of the weights. At this point, VFT dictates that the decision makers preferences be evaluated in terms of weights on the individual attributes in the hierarchy. A basic assumption of this research is that the weights elicited are not constant. It further assumes that a distribution of weights can be developed. While the exact technique for building the distribution from the decision maker's preferences is not considered here, the weight elicitation method itself is important.

3.1 Simulation

For experiments and analysis involving Monte Carlo simulation, certain aspects must be decided before the simulations can be run. First, the various input distributions must be determined. Then, the number of replications needed to achieve the desired research goals must be determined. A random seed is chosen and then random numbers are then drawn from the input distributions and used in accordance with the underlying purpose of the simulation. In this research each set of random draws is used to calculate the weights in Equation 1. The exact mechanics used for designing the simulations in this research are discussed at length.

Once the model has been built, the number of replications needed must be decided. Current computing power makes a very large number of runs both inexpensive and quick. This allows the decision analyst to determine the necessary number of

replications to achieve a desired power or accepted error without being constrained by time or money. The equation for finding sample size is based on a specified Type I and Type II error, or α and β respectively. This allows n to be computed based on desired power and confidence (Hines and Montgomery, 1990: 299). Equation 3 gives the basic form of the equation that was used to calculate the number of replications:

$$n = \frac{(Z_{\alpha/2} + Z_{\beta})^2 \sigma^2}{\delta^2} \quad (3)$$

where $Z_{\alpha/2}$ and Z_{β} come from the cumulative normal distribution for $\alpha/2$ and β respectively, σ^2 = the population variance, and δ^2 = the target difference being detected. While the population variance may not be known, it can be estimated. After each alternative is scored and the single dimension values calculated. Each alternative will exhibit a variance (possibly 0) in its single dimensional values. This variance acts as an upper limit on the final variance of the score for the attribute. Using the largest variance from the group of alternatives can serve as a conservative estimate of the population variance. The variance of a constant times a random variable is equal to the constant squared times the variance of the random variable (Freund and Walpole, 1987: 166-167). The single dimensional value of each attribute acts as the constant and the weight is a random variable. Since all of these values are less than one, their squares will be less than the original value. In essence, the single dimensional values (between 0 and 1) serve to scale down the variance of each weight. When these are aggregated, the sum is likewise scaled down.

3.2 Direct Weighting

In the presence of ambiguity, direct weighting can be used to elicit a distribution or range of possible weights rather than a single constant weight. As was discussed in Chapter 2, this possibility of eliciting a weight distribution is assumed. The methodology being proposed addresses how to handle the distributions once they have been obtained. If this direct weight distribution is used, the weights obtained from the decision maker must follow a logically consistent meaning. Specifically, any possible weight value in the provided distributions must actually be possible within the restriction that the weights sum to 1. If each weight in the specified range of the distribution is not possible, then the decision maker has said that a given attribute could possibly have a weight that it logically could not. This is the inevitable problem with combining independently elicited weight distributions in a fundamentally dependent situation. Each weight must equal one minus the sum of all the other weights. Three possible approaches to resolving this problem are proposed in this section: direct sum, filtered sum, and normalization.

Consider an example using the vehicle decision to illustrate this point. A decision maker weights the hierarchy in the following manner, Cost $\sim U(0.0,0.6)$, Functionality $\sim U(0.4,0.6)$ and Aesthetics $\sim U(0.1,0.3)$. The proposed methodology states that for logical consistency the decision maker must ensure that the most likely values (or means in the case of uniform distributions) sum to 1. Were the decision maker confined to provide only a constant weight, it is assumed that the most likely value would be provided. For this example, $0.3 + 0.5 + 0.2 = 1$, and the test for consistency is passed. The methodology also proposes that every possible value of each distribution must be

realizable. It can immediately be seen that not every possible value of cost can be achieved. If some random draw from the cost distribution equals 0.05, then the sum of the remaining weights must be 0.95 to have the total weights continue to sum to 1. The maximum possible values of the remaining two distributions only sum to 0.9. It is clear that any possible draw of random weights from the provided distributions that sums to 1 will never allow the weight for cost to go below 0.1. Therefore, the provided weights do not allow for every distribution to be completely sampled and the decision maker must be re-engaged to provide weight distributions that are more consistent. Once consistency has been ensured, the method employed must ensure input integrity.

There may be a difference between the decision maker specified weight distribution, the input distribution, and the actual distribution of the weights used in the calculations of the final value. This is an extension of the idea just discussed. If the entire distribution is not sampled, then the actual distribution of the weights used to calculate the final value is not the distribution specified by the decision maker. It is possible for the same loss of input distribution integrity to occur even if each distribution is completely sampled. If this happens, the actual weights used to compute the value of an alternative differ from the distribution the decision maker decided the weight should come from. This disconnect creates a disconnect between the decision maker's true preferences and the actual numbers used to calculate the value of an alternative.

The vehicle example will continue to be used to illustrate the three techniques for direct weighting, but with a new set of weight distributions. Since this is a notional example, $n = 1000$ was chosen. Figure 3 shows the weight distributions that were chosen to illustrate the three direct weighting techniques.




Figure 3: Notional Car Example, with weights

For this example, it is not necessary to identify any alternatives or individual alternative scores. Simply using the weighted hierarchy is sufficient to demonstrate the problems identified. The resulting value from the single dimension value functions would only serve to scale the various elements of the final alternative value.

The first method employed was to simply allow the weights to be randomly generated and summed. This provides a distribution around the final sum as seen in Figure 4.




Figure 4: Direct Weighting, sum of weights

The result from simulating the weights alone with no single dimensional values from the attributes is the same as simulating an alternative who scored perfectly on all measures. The low score on this chart is 0.61. If an alternative were to score perfectly on all measures, its value should be 1.0, not 0.6. The high for this method was 1.33. Again, if 1 is a perfect score, it would be impossible for an alternative to score above that.

Allowing for all possible independent random draws, the direct sum technique, provides results that are inconsistent with reality. It is possible with this technique to have an alternative score almost perfectly and yet still have a low value simply because of the choice of weights. This is the least preferred of the three techniques and should not be used.

The second method for simulating direct weight distributions is to filter out any random pull that is outside a given band. To demonstrate this technique, [0.99,1.01] was chosen. One problem of filtering is that it greatly increases the number of replications needed to get the proper (as calculated with Equation 3) sample size. To demonstrate the results from filtering, the number of replications was increased to 10,000. However, only 373 trials were accepted by the filter. This is a simple hierarchy. When the problems become more complex, even fewer replications will randomly fall within the desired filter. The Figure 5 shows the adjusted value hierarchy for this filtering example.




Figure 5: Filtered Car Value Hierarchy

Closer examination of the weight distributions used show that they meet the range consistency criteria specified previously in this section, however, they lack input distribution integrity. Figure 6 shows the decision maker provided weight distribution for cost as a sloped line representing the triangle distribution specified. It also overlays the actual empirical distribution that was created by the filter and ultimately used in the calculation of alternative values.




Figure 6: Cost weight for filtered car example

An inspection of Figure 6 shows that the distribution used does not actually equal the one specified. This problem appears to be mitigated as input distributions become

less skewed. However, there is no guarantee that a complex decision problem will not contain some fairly skewed distributions. Figure 7 implies that filtering provides a more accurate representation of input distributions when weights are more symmetrically distributed.




Figure 7: Cost weight, symmetric weights

Filtering partially alleviates the problem of the sum of the weights differing greatly from 1.0 by restricting the resulting sum to a narrow interval around 1.0. The cost of this, however is more replications. In the first filtering example, less than 4% of the trials were acceptable. As hierarchies become more complex, this percentage may be even lower. The second problem associated with filtering relates to the integrity of the input distribution. Although all values of the input distributions are possible, the resulting form of the distribution after filtering may not be the same as the input.

The final method available is normalization. For this example, the number of replications was set back to 1000 and the original value hierarchy restored. This method

simply made a random pull from each specified distribution and normalized to 1. The first problem with this is that the interpretation of the weights becomes lost. The process of normalization breaks the connection between the weight and the whole (which should always be 1). When a decision maker specifies a direct weight range from 0.0 to 0.4, the resulting random pull should have the interpretation of the original intent. That is, if the random pulls ends up 0.3, this should be 30% of the final value. Normalizing breaks this. Now, a random draw of 0.3 for Aesthetics could range from 22.2% ($0.3 \div (0.35 + 0.7 + 0.3)$) to 35.3% ($0.3 \div (0.25 + 0.3 + 0.3)$). This may not seem critical, after all it is still in the specified range. However, Figure 8 gives a more glaring example of the problem with normalization from the same value hierarchy .




Figure 8: Cost Weight for Normalized Car Example

Normalization solves the biggest problem of the direct sum technique by simply dividing each attribute weight by the sum of the weights for each random draw. It also avoids the increase in replications caused by filtering. However, as Figure 8

demonstrates, this process ultimately allows for weights that the decision maker did not specify as possible. Normalization can also cause input integrity problems.

This section has demonstrated that while structuring a decision problem with weight ambiguity by using direct weighting may not be too difficult, the actual simulation of that results can be problematic. These examples illustrate a number of problems inherent in each of the proposed techniques. However, if the analyst can afford the time and resources for more replications and if the decision maker has provided symmetrical weight distributions, filtering offers the least deviation from the stated constraints of consistency and integrity.

3.3 Independent Scale Weighting

Since the weights are calculated rather than elicited when using independent scale weighting and swing weighting, the issues addressed in the previous section do not apply. The entire distribution will always be sampled and the process of normalization always ensures that the sum of the calculated weights is 1.

The PWE method as described by Lavelle *et al.* uses a weighting technique based on an independent scale. In this case the scale is the importance to the decision maker and ranges from 0 to 100. Each attribute was independently evaluated against this scale to develop an importance. These importance weights were normalized to provide a weight from 0 to 1. The simulation process for this example did not address the weighting method nor the technique for developing the distributions for the model. Rather, the simulation process made a random draw from each of the input (independent scale weighting) distributions. At each set of random draws, the independent weight

coefficients were normalized. These normalized weights were used to calculate the final values found in the final value distributions. This process is illustrated in the first example in Chapter 4.

3.4 Swing Weighting

As previously noted, swing weighting develops relative relationships between the attributes. Any method for handling ambiguity must preserve this quality. It is inconsistent to develop a set of relative multipliers, normalize them into weights, then try to assess a distribution around these weights. The decision maker has provided his or her preferences on the attributes relative to other attributes. Trying to then obtain distributions based on the numerical weights breaks this connection between the weights of the attributes. Consistent interpretation of uncertain quantities must be maintained. If one attribute is twice as important as another in the most likely case, all other possible relationships must be made under the same condition of one attribute relative to another. To preserve this, the swing weight coefficients are the quantities that are allowed to vary.

When weights are obtained through swing weighting, many of the issues discussed previously are no longer applicable and some of the difficulties encountered with direct weighting are similarly resolved. Because swing weights are defined relative to each other, they must be normalized in order to develop weights that can be appropriately used in Equation 1. Because of this inherent need to normalize, requiring that any given random draw sums to one is moot. Similarly, the requirement that the input distribution be fully sampled is also no longer a factor. Since the specified

distributions are sampled before creating weights that sum to 1, every part of the input distribution can always be used, no matter how small or large. Normalization will simply force the other weights to change to compensate for this. Finally, input integrity is no longer as important. The numerical weights used in calculations provide only indirect resemblance to the swing weight coefficients provided by the decision maker. The distribution of these numerical weights will likewise have little connection (in terms of distribution form or parameters) to the input distribution.

The commonality between swing weighting and direct weighting is that the form of the input distributions will impact whether the mean of the final alternative distributions equal the result of using the most likely case as point estimates. This issue is not, however, as clear cut as it is with direct weighting. In swing weighting, an input distribution can be symmetric, but the normalization process calculates weights in a disproportionate manner.

Standard swing weighting identifies the lowest ranked attribute and holds this as the baseline. As long as the baseline stays constant, any swing weight coefficient (not necessarily the lowest) can be used.

3.5 Selecting Alternatives in the Presence of Ambiguity

When Value-Focused Thinking is used to establish single constant values for alternatives in question, ranking the alternatives is a fairly straightforward matter. Human judgment and sensitivity analysis can be used to provide insight into alternatives whose overall values are relatively close. These alternatives can be evaluated further, but

the results will again be some point estimate of value that can be easily ordered.

Selecting the best alternative is not, however, as simple when the values of the alternatives are represented by some probability distribution.

3.5.1 Dominance

Dominance has been used when uncertainty in the alternatives' scores are present as a means of identifying alternatives that are more desirable than others (Eum *et al.*, 2001; Lee *et al.*, 2001; Howard, 1966: 100-102; Langweich and Choobineh, 1996).

There are two forms of dominance to consider: deterministic and stochastic. In the condition where more is preferred, deterministic dominance occurs when the lowest possible value of an alternative (A) is higher than the highest possible value of another alternative (B) (Clemen, 1996: 123-127). In this case we say that alternative A deterministically dominates alternative B. Practically speaking, if a decision maker can always get a higher value with A, regardless of variation, there is little reason to select B.

The second form of dominance is stochastic dominance (Clemen, 1996: 123-127).

Stochastic dominance can most readily be seen by plotting the cumulative distribution function (CDF) of the alternatives (Clemen, 1996: 123-127). If alternative A has an equal or greater probability than alternative B at every possible value (from 0 to 1), we say A stochastically dominates B. Figure 9 demonstrates stochastic dominance.




Figure 9: Stochastic Dominance

In this example, the alternative labeled ESM stochastically dominates the Baseline. The Baseline starts at a lower overall value and at no point in the cumulative distribution is more of the baseline at or above a higher overall value.

3.5.2 Statistical Tests

One of the advantages of having an empirical distribution is that it allows for statistical analysis on the means. For any two given alternatives, a hypothesis test can be used to determine if there is a significant difference in the means or variances. It may also be appropriate if the standard error of the mean is “sufficiently large.”

Unfortunately, “sufficiently large” must be judged on a case-by-case basis.

The following hypothesis test can be used for identifying statistical differences in the means for any two alternatives (Hines and Montgomery, 1990: 312-315). To follow

the previous example: $\mu_D = \frac{\sum_{j=1}^n V_{ESM} - V_{Baseline}}{n}$ and the resulting hypothesis test would

be:

$$\begin{aligned} H_0 : \mu_D &= 0 \\ H_a : \mu_D &\neq 0 \end{aligned} \tag{4}$$

This is a paired t-test. Since for any given random draw of weights, two alternatives are linked by those given weights, a paired t-test is applicable. In the test outlined here, the null hypothesis, H_0 , says that the difference in means is zero. In this case, the test is attempting to discern if the means are equal. The test statistic is:

$$t_0 = \frac{\bar{D}}{S_D / \sqrt{n}} \quad \text{where} \quad \bar{D} = \frac{\sum_{j=1}^n D_j}{n} \quad \text{and} \quad S_D^2 = \frac{\sum_{j=1}^n D_j^2 - \left[\left(\sum_{j=1}^n D_j \right)^2 / n \right]}{n-1} \tag{5}$$

where:

$\bar{D} \equiv$ the overall mean difference

$n \equiv$ the number of replications

$S_D^2 \equiv$ the sample variance of the differences

In this case, $D_j = X_{1j} - X_{2j}$, $j = 1, \dots, n$, the difference between the alternative values at any given draw. The null hypothesis is rejected if $t_0 > t_{\alpha/2, n-1}$ or if $t_0 < -t_{\alpha/2, n-1}$.

3.6 Methodology Summary

After the decision hierarchy has been built, the weight elicitation method must be chosen and the simulation model built. Once the model has been built, an appropriate n must be determined. Following the determination of n , the exact process for simulating the problem is dependent on the weight elicitation method chosen. If direct weighting is used, a number of issues are involved that must be considered to maintain consistent and meaningful interpretation of the weights. There are a number of different techniques to simulate these direct weights. While the filter method still has shortcomings, this approach is by far the preferred method. Independent scale weighting and swing weighting are very similar in that they both resolve many of the issues involved in direct weighting. They are also similar in their use of normalization to transform decision maker preferences into weights. In both independent scale weighting and swing weighting, the information simulated is not the weights, but rather those relative quantities provided by the decision maker. Figure 10 maps the methodology presented in this chapter.

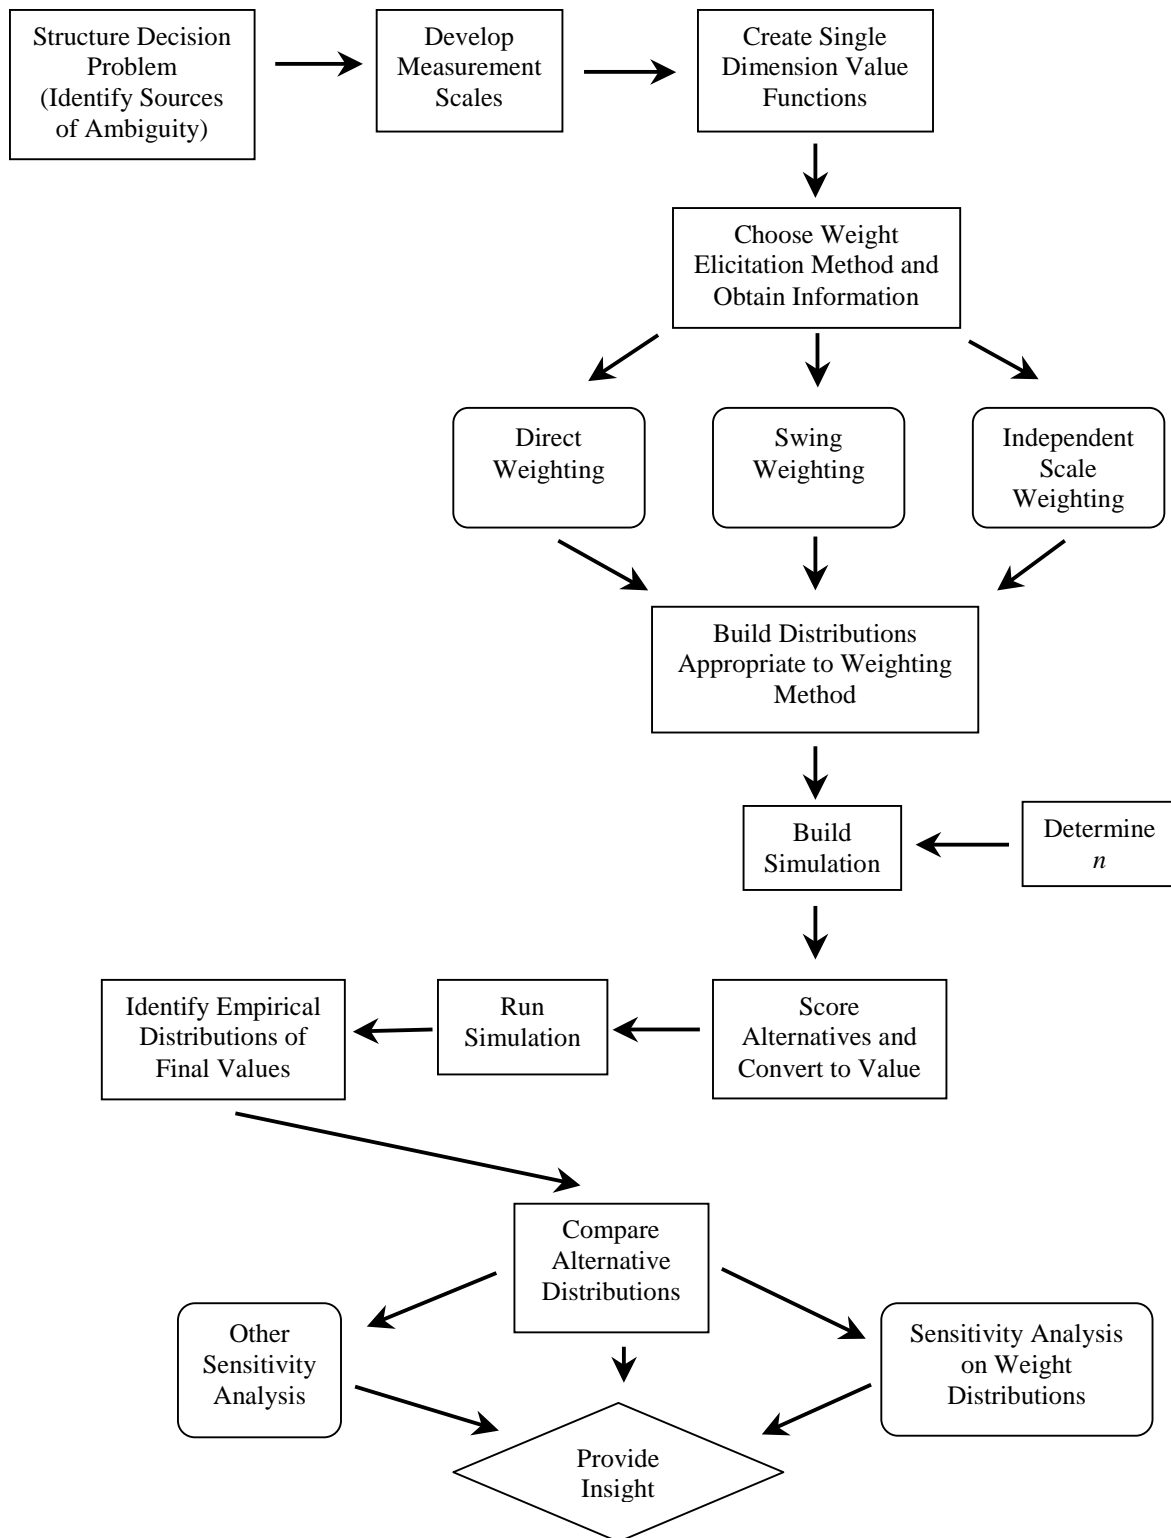


Figure 10: Methodology Flow Chart

The methodology presented in this chapter follows the flow chart provided in Figure 10. The steps given are further described as: 1) Identify and structure the objectives and criteria important to a decision maker for the decision situation at hand. This also includes identifying any sources of ambiguity. Specifically, identify if there will be any ambiguity in the weights obtained later in the process. 2) Develop measurement scales for each of the final end measures. 3) Create Single Dimension Value Functions to convert raw measurement scores to values. 4) Using the decision structure and possible sources of ambiguity, determine which weight elicitation method (direct, swing or independent scale) is most appropriate. 5) Using the weight elicitation method chosen, obtain relevant weight information from the decision maker. 6) With the information just obtained, build the input distributions needed to determine weights. 7) Determine the appropriate number of replications needed. 8) Build the simulation based on the elicitation method chosen and the distributions available. 9) Score each of the identified alternatives on the measurement scales and convert them to values. 10) Run the simulation for the number of replications determined. 11) Identify the empirical distributions, and their various parameters (e.g. mean and variance), for each alternative that result from the simulation. 12) Compare alternative distributions through observation, dominance and statistical tests on parameters. 13) Perform sensitivity analysis on the form and parameters of the input distributions. Other sensitivity analysis, as appropriate, can also be conducted. 14) Through the comparisons in step 12 and the sensitivity analysis in step 13, provide insight to the decision maker on the relative value of the alternatives and any possible consequence of realistic changes in the decision structure or inputs.

This flow chart does not represent a single process, but rather an iterative approach in which many steps may result in reevaluating some previous step. The possible iterations and feedback loops have been omitted from the flow chart to eliminate confusion and highlight the primary flow of the process.

This chapter opened with a discussion of building a simulation model. The three weight elicitation methods under consideration were then discussed. This discussion included some of the specific techniques that could be employed to incorporate the weight distributions into the simulation model. Finally, some possible methods for evaluating the resulting value distributions were outlined. These ideas lead directly into Chapter 4 where two simulation examples are presented and their results used to strengthen the proposed methodology in this thesis.

4. Analysis

Chapter 3 introduced a methodology to simulate the effects of weight ambiguity in a Value-Focused Thinking decision problem. It also provided the ground work for modeling VFT problems with weight ambiguity. Chapter 2 established the validity of the Monte Carlo approach and this chapter provides concrete examples of the methodology.

Section 4.1 provides the simulation results of the Airport Siting Problem used by Lavelle *et al.* Statistical hypothesis testing is then conducted on the mean and variance of the empirical distributions to compare them to the first and second moments, respectively, of the analytical results. Section 4.2 focuses on the results of using weight ambiguity on a more complex decision problem. Beauregard's model for Information Assurance is analyzed using notional distributions in place of constant weights. The resulting value distributions for each alternative are presented followed by a discussion of several ways to compare the alternatives to provide insight. The chapter concludes with a brief summary of the results.

4.1 Airport Citing

The importance of simulating the Airport Siting example from Lavelle *et al.* is to provide a link between the analytical results and the simulated results. As is shown in this section, the simulation results approximate the analytical results very closely. This

result is used later as a first look at the accuracy of the simulation methodology when used on more complex decision problem structures.

When possible, it is generally considered better practice to solve a problem analytically rather than through the use of simulation. Lavelle *et al.* used a greatly simplified decision structure to develop an analytical solution to the problem of parametric uncertainty. As the structure of the problem becomes more complex, these analytical methods become cumbersome and may even be intractable (Hammersly and Handscomb, 1964). In these cases, simulation becomes increasingly important. Simulation can often be employed to provide greater speed and flexibility to the decision analyst. By simulating the simplified structure used by Lavelle *et al.*, this section provides a benchmark between the simulated and analytical results, highlighting the approach while illustrating it with an equivalent outcome.

4.1.1 Building Model

To demonstrate the ability of Monte Carlo Simulation to duplicate the analytical results from Lavelle *et al.*, the simulation model was constructed as near identical as possible to the problem structure used in the original study. The following information was provided as the weight distribution information:

Table 1: Attribute weight data, from Lavelle *et al.*, 1997: 779

Attribute	Minimum	Mode	Maximum	Distribution
1. Cost	40.00	--	60.00	Uniform
2. Capacity	30.00	90.00	95.00	Triangular
3. Access Time	30.00	85.00	88.00	Beta
4. Safety	15.00	--	80.00	Uniform
5. Displacement	28.00	85.00	92.00	Beta
6. Noise	20.00	--	100.0	Uniform

The following three tables give the uncertainty data for the ratings of the alternatives. For each alternative, the uncertainty distribution information for each of the three attributes is given.

Table 2: Rating data for Alternative A, from Lavelle *et al.*, 1997: 780

Alternative A				
Attribute	Minimum	Mode	Maximum	Distribution
1. Cost	50.00	65.00	84.00	Triangular
2. Capacity	50.00	68.00	95.00	Beta
3. Access Time	50.00	--	78.00	Uniform
4. Safety	61.00	75.00	85.00	Beta
5. Displacement	60.00	75.00	79.00	Triangular
6. Noise	30.00	68.00	89.00	Beta

Table 3: Rating data for Alternative B, from Lavelle *et al.*, 1997: 780

Alternative B				
Attribute	Minimum	Mode	Maximum	Distribution
1. Cost	62.00	65.00	68.00	Triangular
2. Capacity	63.00	66.00	80.00	Beta
3. Access Time	55.00	--	70.00	Uniform
4. Safety	65.00	75.00	80.00	Beta
5. Displacement	65.00	70.00	77.00	Triangular
6. Noise	55.00	65.00	70.00	Beta

Table 4: Rating data for Alternative C, from Lavelle *et al.*, 1997: 780

Alternative C				
Attribute	Minimum	Mode	Maximum	Distribution
1. Cost	33.00	88.00	95.00	Triangular
2. Capacity	25.00	88.00	92.00	Beta
3. Access Time	7.00	--	100.0	Uniform
4. Safety	14.00	88.00	89.00	Beta
5. Displacement	30.00	85.00	91.00	Triangular
6. Noise	35.00	84.00	99.00	Beta

The information contained in Tables 1-4 was used to construct a Monte Carlo Simulation model. Table 1 provided the weight distributions. Tables 2-4 provided the single dimensional values for each of the six attributes. Because Lavelle *et al.* used independent scale weighting, each random draw from the distributions provided in Table 1 was normalized to 1. For each attribute the normalized weight was multiplied by the single dimensional value and the six resulting products were summed to give a final value. This was done with each of the three alternatives.

4.1.2 Determining n

Following the construction of the simulation model, Equation 3 was used to determine the number of replication required. To determine the number of replications using Equation 3, α and β were both chosen as 0.05 with a resulting 5% chance of making either a Type I or Type II error for any given δ . 0.05 was chosen arbitrarily as a commonly accepted standard. The Lavelle *et al.* study scaled all weights and single dimensional values up two orders of magnitude and the resulting calculations reflect this. In Equation 3, δ represents the smallest difference in means that can be detected with the

given power. The smallest difference in means of the alternatives given by the analytical solution in Lavelle *et al.* is 1.12. Since this is the known difference from the analytical solution, choosing $\delta = 1$ would indicate that if the distributions of the alternatives are compared using the simulation results, the same level of difference among them would be detectable. The given analytical results also provided a maximum variance in the alternative scores of 51.41. As was discussed in Chapter 3, this variance can be used as an estimate for the population variance used to calculate n .

$$n = \frac{\left(Z_{\alpha/2} + Z_{\beta}\right)^2 \sigma^2}{\delta^2} = \frac{(1.96 + 1.645)^2 \cdot 51.41}{1} = 668.1 \approx 669$$

4.1.3 Simulation Results

The following table and figures provide the results from the simulation both numerically and graphically. Table 5 gives the mean and variance from both the analytical and empirical distributions for each alternative.

Table 5: Simulating Lavelle

Alternative	PWE		Simulation	
	Mean	Variance	Mean	Variance
A	68.32	9.25	68.34	10.46
B	67.21	1.34	67.22	1.76
C	69.88	51.41	69.27	58.17




Figure 11: Alternative A

Figure 11 provides the final distribution of the value of Alternative A. In Figure 11, the vertical bars come from the numerical results of the simulation and are labeled as “Alternative A” in the legend. The solid line in the figure is the best fit normal distribution for the given data. The legend provides the mean and standard deviation of this best fit curve. This figure, as well as Figures 12 and 13, seems to indicate by inspection that the simulated and theoretical results are quite close. Exactly how close will be explored in the following section.




Figure 12: Alternative B

Figure 12 provides the final distribution of the value of Alternative B. The bars and solid line are interpreted in the same way as they were in Figure 11, the numerical results and the best fit curve.




Figure 13: Alternative C

Figure 13 provides the final distribution of the value of Alternative C. As with Figures 11 and 12, this final chart gives the numeric results as bars and the best fit curve as a solid line.

4.1.4 Comparing Results

The key to this portion of the research is to demonstrate that the simulated results closely approximate the analytical results, thereby allowing for a method that does not require cumbersome integration. The validation of simulation as an approach to weight ambiguity in the simplified problem structure of PWE also provides a foundation to begin to simulate more complex problems.

To demonstrate that the simulation results closely approximate the analytical results, two sets of hypothesis tests were conducted. The first set is designed to compare the means of the analytical and empirical distributions and the second is designed to test the variance. Equation 6 provides these tests:

$$\begin{array}{ll} H_0 : \mu = \mu_0 & H_0 : \sigma^2 = \sigma_0^2 \\ H_a : \mu \neq \mu_0 & H_a : \sigma^2 \neq \sigma_0^2 \end{array} \quad (6)$$

where μ_0 and σ_0^2 represent the analytical mean and variance respectively and μ and σ^2 represent the mean and variance calculated from the empirical distributions.

In the case of the means, the test statistic used is (Hines and Montgomery, 1990: 301):

$$Z_0 = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}} \quad (7)$$

where μ_0 and \bar{X} are the analytical mean and empirical mean respectively, σ is the standard deviation from the analytical solution and n is the number of replications already determined. In this case, if Z_0 falls between $-Z_{\alpha/2}$ and $Z_{\alpha/2}$ then we are unable to reject the null hypothesis that the analytical and empirical means are equal. The same confidence level, $\alpha = 0.01$, is being used as it has been previously.

Comparison of the variance will use the test statistic (Hines and Montgomery, 1990: 317):

$$Z_0 = \frac{S - \sigma_0}{\sigma_0/\sqrt{2n}} \quad (8)$$

where σ_0 and S are the analytical standard deviation and the empirical standard deviation respectively and n is the number of replications already determined. In this case, if Z_0 falls between $-Z_{\alpha/2}$ and $Z_{\alpha/2}$ then we are unable to reject the null hypothesis that the analytical and empirical standard deviations, and consequently the variances, are equal. The same confidence level, $\alpha = 0.01$, is being used as it has been previously. This test statistic was chosen as a large-sample test that is robust to errors in the normality assumption (Hines and Montgomery, 1990: 317-318).

The first test for the mean. For Alternative A:

$$Z_A = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}} = \frac{68.34 - 68.32}{3.04/25.87} = \frac{0.02}{0.118} = 0.17$$

For Alternative B:

$$Z_B = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}} = \frac{67.22 - 67.21}{1.16/25.87} = \frac{0.01}{0.045} = 0.22$$

For Alternative C:

$$Z_C = \frac{\bar{X} - \mu_0}{\sigma/\sqrt{n}} = \frac{69.27 - 69.88}{7.17/25.87} = \frac{-0.61}{0.277} = -2.20$$

For all three alternatives, the critical values and rejection regions are the same.

With $\alpha = 0.01$, the critical values given by $-Z_{\alpha/2}$ and $Z_{\alpha/2}$ are -2.58 and 2.58 respectively.

For Alternatives A, B and C, it is clear that the test statistic is within this range and therefore the null hypothesis cannot be rejected. This indicates that the simulated results are statistically equivalent to the analytical results.

After the means are compared, the variances are tested. Again, the hypothesis test and test statistic were both given previously in this section. The results are, for

Alternative A:

$$Z_A = \frac{S - \sigma_0}{\sigma_0/\sqrt{2n}} = \frac{3.23 - 3.04}{3.04/36.58} = \frac{0.19}{0.083} = 2.29$$

For Alternative B:

$$Z_B = \frac{S - \sigma_0}{\sigma_0/\sqrt{2n}} = \frac{1.33 - 1.16}{1.16/36.58} = \frac{0.17}{0.032} = 5.31$$

For Alternative C:

$$Z_C = \frac{S - \sigma_0}{\sigma_0/\sqrt{2n}} = \frac{7.63 - 7.17}{7.17/36.58} = \frac{0.46}{0.196} = 2.35$$

The critical values and rejection regions are the same for all three alternatives. With $\alpha = 0.01$, the critical values given by $-Z_{\alpha/2}$ and $Z_{\alpha/2}$ are -2.58 and 2.58 respectively. Comparing these critical values to the computed test statistics results in being unable to reject the null hypothesis that the simulated and theoretical variances are equal for Alternatives A and C. For Alternative B, the null hypothesis, that the simulated and theoretical variances are equal, would be rejected. This indicates that for Alternative B, there is a significant difference in the simulated variance and the theoretical variance.

The purpose of this example was to demonstrate that the results from the simulation would approximate the results from the analytical solution. Five out of the six tests performed failed to reject the null hypothesis. That is, for five tests there was not enough evidence to conclude that the simulated and analytical moments being compared were different.

4.2 Information Assurance

Simulating Information Assurance provides a look at the results of the proposed method on a more complex, swing weighted value hierarchy. None of the assumptions necessary to solve a decision problem with weight ambiguity analytically apply to this problem. The only thing that has been changed in this problem from the original is in substituting the constant weights with distributions.

4.2.1 Building the Model

The previous section has shown that simulating the input parameters to a weighted additive model provides results equivalent to the analytical solution. While this may not be necessary when simple problem structures such as PWE are used, the situation changes considerably when complex decision structures are modeled. In these cases, the analytical solution can quickly become time consuming and difficult, if not intractable.

A hierarchically structured VFT problem using swing weighting is just such a case. Not only do these structures potentially contain the distributions used by Lavelle *et al.*, they may also be specified empirically or through some mixed distribution. The hierarchical nature of the VFT methodology would now cause the final weights to be not simply normalized, but also a product of several local weights.

The decision hierarchy used in Figure 2 forms the basis of a more complex decision structure which now incorporates weight ambiguity. One possible source of ambiguity in this decision problem is designing a hierarchy that can be used not only in the organization it was built for, but to model Information Assurance problems for other organizations. In such a case, the preference tradeoffs of several decision makers need to be taken into account. Figure 14 contains the information from Figure 2, but also includes the corresponding distribution for each of the swing weight coefficients. In each case, the original provided swing weight coefficient will represent the most likely value for each distribution. This may be different from the expected value of the distribution.

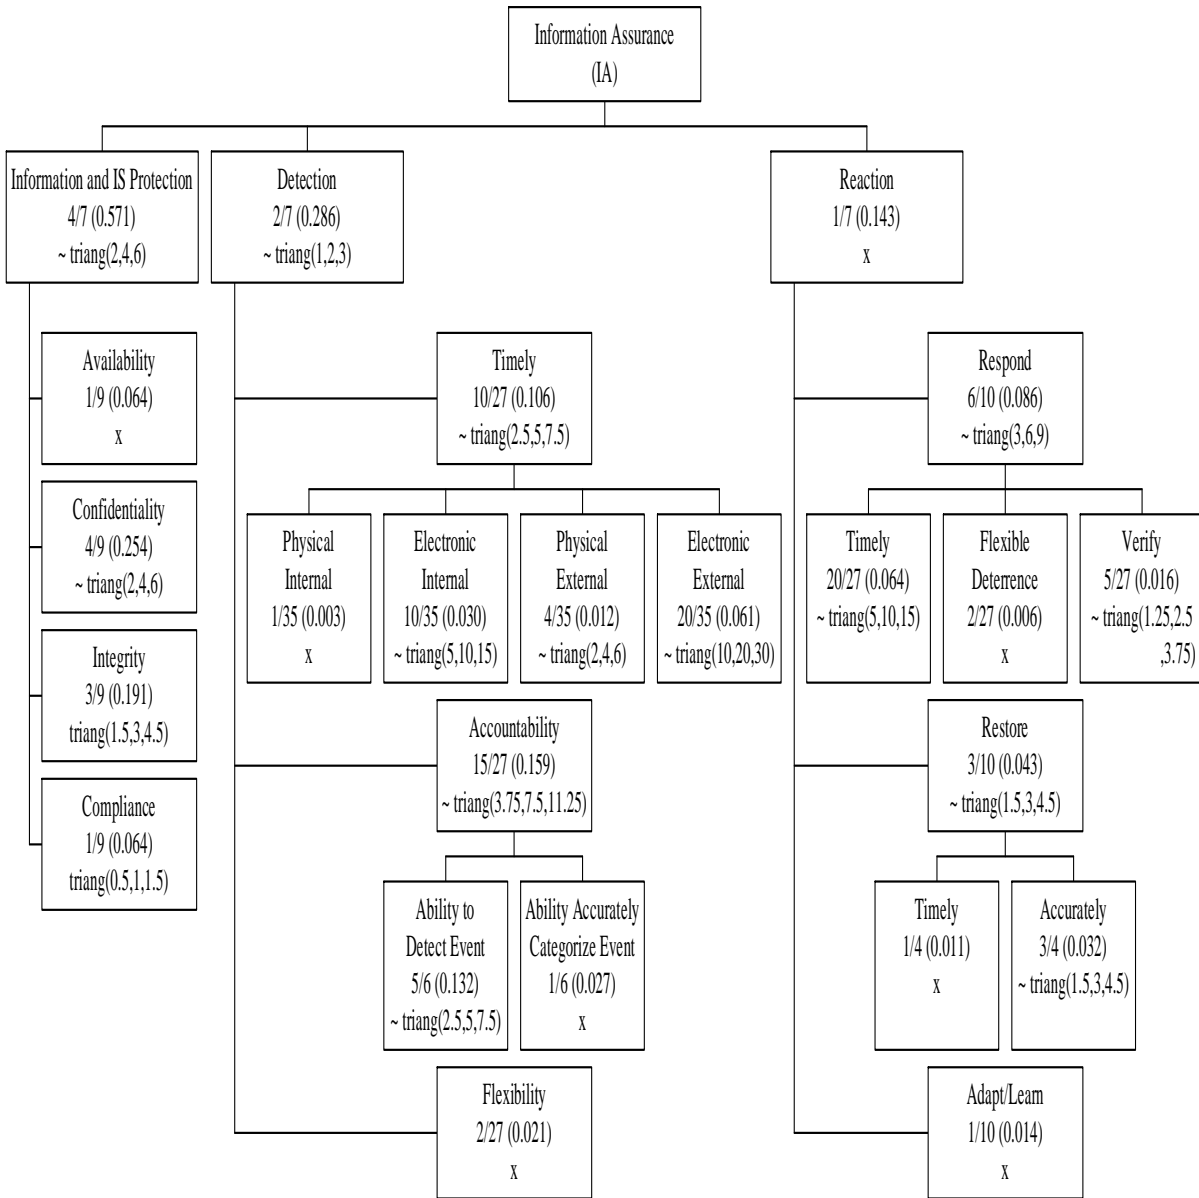


Figure 14: IA Hierarchy with Swing Weight Coefficient Distributions

With no clear guidance from the original study on how to structure the weight distributions, it was decided to simply provide all weights with a triangular distribution with the original swing weight coefficient as the mode, half this value for the min and one and a half times this value for the max. This created all symmetrical distributions.

There is one case that should be noted. Under Information and IS Protection, two attributes were deemed to have equal weight and also be the least important. One of them was arbitrarily chosen to remain equal to x , thereby representing the least important attribute. The other was distributed $\sim \text{triang}(0.5, 1, 1.5)$.

As described in Chapter 2, swing weights are developed by identifying the least important attribute and using it as the baseline, x . Each successive attribute in the given tier and branch is then compared to this attribute and its swing weight coefficient is equal to its relative importance above the baseline. Since all swing weight coefficients and corresponding local weights are identified against the value of the baseline, this value must not change. As shown in Figure 14, the coefficient of the swing weight distribution for the least important attribute is fixed at 1. However, since the weights are normalized after each random draw, the actual distribution of the corresponding numerical weight will not be a constant.

4.2.2 Determining n

As was done for the Lavelle *et al.* study, n must be determined for this simulation. The number of replications was again calculated from the desired Type I error, Type II error and detectable difference. Equation 3 was again used. In this case, the hypothesis test was conducted on the distributions of the final values of the alternatives. Without any prior knowledge of the final alternative distributions, the population variances are not known. However, there is information available that can help provide an estimate. As

was described in Chapter 3, the population variance can be estimated by using the variance of the single dimensional values for the final end measures for each attribute.

For the sample size calculations for this example, $\alpha = \beta = 0.05$. This provides a typical 95% confidence, and a power of 0.95 for any given δ . The choice of δ is based on a desire to be able to detect a certain difference. In this case, it is desirable to detect a minimum difference of 0.01. The largest variance in the component utilities of the alternatives is used as an approximation of the variance. In the Information Assurance example, the largest variance is 0.0039. This is the variance used to calculate n .

$$n = \frac{(Z_{\alpha/2} + Z_{\beta})^2 \sigma^2}{\delta^2} = \frac{(1.96 + 1.645)^2 \cdot 0.0039}{0.01^2} = \frac{0.05}{0.01^2} = 500$$

4.2.3 Simulation Results

Table 6 provides a summary of parameters associated with the empirical distributions for the final value of the alternatives. For reference it also includes the original value of each alternative.

Table 6: Information Assurance Simulation Results

Statistic	Alternatives			
	Baseline	ISS	ESM	Cisco
Original Value	0.6180	0.7530	0.6830	0.9100
Mean	0.6190	0.7532	0.6824	0.9095
Variance	0.0001	0.0001	0.0002	0.0001
Range Min	0.5876	0.7228	0.6463	0.8736
Range Max	0.6519	0.7758	0.7244	0.9302

The results in Table 6 clearly show that the distribution mean is very close to the point estimate mean and that the variances are quite low. While this table provides much useful information, it does not give an intuitive look at the distributions of the alternatives. Figures 15-18 provide this intuitive look.




Figure 15: IA – Baseline

Figure 15 gives the distribution of the final value of the Baseline system. The chart includes vertical bars representing the exact numerical results from the simulation. The solid line is the best fit normal approximation to the data. The legend provides the mean and standard deviation for this best fit line.




Figure 16: IA – ISS

Figure 16 provides the numerical results from simulating the value of the ISS alternative. Again, a best fit normal approximation is provided along with the mean and standard deviation.




Figure 17: IA – ESM

Figure 17 give the final distribution of value for the ESM Information Assurance alternative. This chart includes the numerical results of the simulation as a series of

vertical bars with a best fit approximation of the normal as a solid line. The mean and standard deviation of the normal is given in the legend.




Figure 18: IA – Cisco

Figure 18 gives the same information as the previous three figures for the Cisco alternative. Again, the mean and standard deviation of the normal approximation is given in the legend.

Figures 15-18 show the empirical distribution and fitted normal distribution for each of the alternative. Figure 19, however, truly provides an idea of how the alternatives compare.




Figure 19: Information Assurance Alternatives

It is clear from Figure 19 that the inclusion of ambiguity in this model, assuming that the ambiguity was properly specified, does not change how this decision would have been made. This figure also indicates that Cisco is by and far the best alternative. For this decision problem, part of the insight provided to the decision maker is that within the specified probable weights, there is not set of weights in which any alternative outperforms Cisco. This information is useful in determining that the Cisco alternative is clearly a better alternative

4.2.4 Dominance

Once the simulation has been built and run, the results can begin to be analyzed. The obvious first method is to simply look at the distributions of the alternatives when plotted together, as shown in Figure 19. This may clearly rule out some alternatives or

separate the alternatives into distinct groups. After that, however, the comparison becomes more difficult.

The first level of analysis is to check for dominance among the alternatives as was discussed in Chapter 3. One alternative (A) deterministically dominates another alternative (B) if the lowest possible value of A is higher than the highest possible value of B. Finding dominance among the alternatives allows the decision maker to draw strong conclusions about which alternative is better.

For the Information Assurance example, dominance can easily be determined from Table 6 by looking at the min and max range values. By comparing these values, Cisco clearly dominates all other alternatives. In addition, ISS dominates the Baseline system. There is no deterministic dominance between ESM and both the Baseline and ISS. Stochastic dominance, however, exists between any two alternatives available. In these cases, it is clear that Cisco displays deterministic dominance over all others. ISS dominates both the Baseline and ESM and finally, ESM dominates the Baseline. These results come as no surprise given the original result. Figure 20 illustrates the presence of stochastic dominance among the four alternatives.




Figure 20: Cumulative distribution chart showing stochastic dominance

Figure 20 clearly shows the stochastic dominance present among the four alternatives. This clearly indicates that the Cisco alternative would be preferred.

Chapter 4 presented two examples used to demonstrate the proposed methodology. The first example, Airport Citing, demonstrated the connection between analytical and simulated results. It provided a strong indication that the proposed methods of simulation accurately model the reality of using probability distributions to represent weight ambiguity in a decision problem. The second example, Information Assurance, illustrated the usefulness of simulation when incorporating weight ambiguity into larger, more complex decision structures. It also provided a venue for briefly discussing possible comparison methods for resulting distributions of value for the alternatives. Together, the examples demonstrate the validity and applicability of simulation.

5. Conclusions and Discussion

With the inclusion of weight ambiguity into the VFT methodology, the decision analyst is able to bring a more flexible and robust process to the decision maker. “[I]t creates a paradigm for *a priori* sensitivity analysis, thereby giving the decision maker more information upon which to base decisions...” (Lavelle *et al.*, 1997: 774). Although the presence of weight ambiguity has been a long established fact, most decision analysis methodologies simply assumed it away. The methodology presented here, and demonstrated by the examples, makes that assumption no longer necessary.

5.1 A New Approach

The technique presented here is an adaptation and generalization of Keeney’s VFT process. It has been expanded to identify and include potential ambiguity in the weights. While the idea that a decision maker’s weights may not be absolute is not new and despite the increasing prevalence of groups in the decision making process, there has been very little research into expanding current methods to account for preference ambiguity.

A new approach utilizing Monte Carlo simulation is clearly suggested. The methodology proposed here begins to fill that void. Comparison with the analytical results from Lavelle *et al.* indicate a very close congruity between the theoretical and simulated results. Given the current state of computing power, a simulation approach offers a greater degree of flexibility to the decision analyst. Once the simulation model has been built, modifying the model for “what if” analysis and sensitivity analysis takes

little time. Furthermore, this procedure is both robust enough to handle decision problems of varying complexity and does so in a manner easy to duplicate so it may be put into practical application, provided the required distributions can be elicited.

5.2 Results

The hypothesis tests resulting from Airport Siting example suggest the applicability of Monte Carlo simulation. Five of the six tests showed no statistical difference between the analytical solution and the simulated results. The statistical results and the accompanying figures clearly, although not perfectly, demonstrate that simulation can closely approximate an analytical solution.

The Information Assurance decision problem was represented by a complex and multi-tiered structure. This was also successfully simulated. While the results did not lend themselves to a more descriptive assessment of the different alternative selection options, they showed that distributions can be developed. Even this lack of alternative overlap is insightful. The lack of substantial overlap illustrated both deterministic and stochastic dominance. Identifying the presence of dominance in a decision problem is also insightful. This may help identify the clearly preferred alternative even if the decision maker's preference structure varies widely.

The direct weighting discussion and examples shed particular light on the potential shortcomings of direct weighting when applied to weight ambiguity. The fundamental problem arises from using independent random variable to model a dependent reality. The issues themselves demonstrate that the idea of using ambiguity in

weights must be closely scrutinized to assure that the information obtained from the decision maker is not inconsistent.

5.3 Future Work

One of the largest areas for future research is in the area of sensitivity analysis. Since this method directly incorporates changes in weights, is there any need for traditional sensitivity analysis? Are there possibly other types of sensitivity analysis that may be useful? Sensitivity analysis for constant weights centers around varying those weights to determine how robust the decision is to changes. With distributions, it may now be necessary to look at sensitivity analysis in terms of varying the distribution of the weights. The impact on the decision when certain weights are given more or less variance may be used. In addition, traditional sensitivity analysis using constant weights may still be effective if the weight distributions have been incorrectly elicited. There may be ways to incorporate a more deliberate and insightful sensitivity analysis. Additional methods of comparing and assessing the final alternative distributions may open up avenues to more insight into the decisions. Research into the best methods for eliciting weight distributions should be developed. Finally, the model needs to be applied in a number of settings to discover its strengths and try to bring to light any flaws. More work can also be done in the testing of more input distributions to the model. The most significant advances to this research would be in the fields of probability elicitation from decision makers and the comparison of alternatives.

Bibliography

- Beauregard, Joseph E. *Modeling Information Assurance*. MS Thesis, AFIT/GOR/ENS/01M-03. School of Engineering and Management, Air Force Institute of Technology (AU), Wright-Patterson AFB OH, March 2001.
- Benson, George, Shawn Curley and Gerald Smith. "Belief Assessment: An Underdeveloped Phase of Probability Elicitation," *Management Science*, 41 (10): 1639-1653 (October 1995).
- Bernoulli, Daniel. "Exposition of a New Theory on the Measurement of Risk," *Econometrica*, 21: 23-36 (1954). Translated by Dr. Louise Sommer.
- Beroggi, Giampiero. *Decision Modeling in Policy Management – An Introduction to the Analytic Concepts*. Boston: Kluwer Academic Publishers, 1999.
- Bottomley, Paul and John Doyle. "A comparison of three weight elicitation methods: good, better, and best," *Omega: The International Journal of Management Science*, 29: 553-560 (2001).
- Buchanan, John, James Corner and Mordecai Herig. "Assumptions of Multiple Criteria Decision Models," Presented at INFORMS: Miami Beach FL, November 2001.
- Bunn, Derek W. *Applied Decision Analysis*. New York: McGraw-Hill Book Company, 1984.
- Canada, John R. and William G. Sullivan. *Economic and Multiattribute Evaluation of Advanced Manufacturing Systems*. Englewood Cliffs, NJ: Prentice Hall, 1989.
- Chacko, George K. *Decision-Making under Uncertainty: An Applied Statistics Approach*. New York: Praeger, 1991.
- Choo, Eng, Bertram Schoner, and William Wedley. "Interpretation of criteria weights in multicriteria decision making," *Computers and Industrial Engineering*, 37: 527-541 (1999).
- Clemen, Robert T. *Making Hard Decisions: An Introduction to Decision Analysis*, 2nd Ed. Boston MA: Duxbury Press, 1996.
- Davison, A. C. and D. V. Hinkley. *Bootstrap methods and their applications*. Cambridge: Cambridge University Press, 1997.
- Ellis, Donald and B. Fisher. *Small Group Decision Making: Communication and the Group Process*, 4th Ed. New York: McGraw-Hill, Inc., 1974.

- Eum, Yun Seong, Kyung Sam Park and Soung Hie Kim. "Establishing dominance and potential optimality in multi-criteria analysis with imprecise weight and value," *Computers & Operations Research*, 28: 397-409 (2001).
- Felli, James and Gordon Hazen. "Do sensitivity analyses really capture problem sensitivity? An empirical analysis based on information value," *Risk, Decision and Policy*, 4 (2): 79-98 (1999).
- Fishburn, Peter. *Utility Theory for Decision Making*. New York: John Wiley & Sons Inc., 1970.
- Fishburn, Peter. *Decision and Value Theory*. New York: John Wiley & Sons Inc., 1964.
- Freund, John and Ronald Walpole. *Mathematical Statistics, Fourth Edition*. Englewood Cliffs NJ: Prentice-Hall Inc., 1987.
- Frey, H. "Separating Variability and Uncertainty in Exposure Assessment: Motivations and Method," *Proceedings of the 86th Annual Meeting, Air and Waste Management Association*. June, 1993.
- Hamersley, J. M. and D. C. Handscomb. *Monte Carlo Methods*. New York: John Wiley & Sons Inc., 1964.
- Hines, William and Douglas Montgomery. *Probability and Statistics in Engineering and Management Science, Third Edition*. New York: John Wiley & Sons, Inc., 1990.
- Howard, Ronald A. "Decision Analysis: Applied Decision Theory," *Proceedings of the Fourth International Conference on Operational Research*. New York: Wiley-Interscience, 1966: pp 55-71.
- Howard, Ronald A. and J.E. Matheson. "An Introduction to Decision Analysis," *The Principles and Applications of Decision Analysis, Volume 1*. Stanford Research Institute, 1968a: pp 20-55.
- Howard, Ronald A. "The Foundations of Decision Analysis," *IEEE Transactions on Systems Science and Cybernetics*, SSC-4 (3): 1-9 (September 1968b).
- Howard, Ronald A. "An Assessment of Decision Analysis," *Operations Research*, 28 (1): 179-202 (January-February 1980).
- Howard, Ronald A. "The Evolution of Decision Analysis," in *The Principles and Applications of Decision Analysis, Volume 1*. Ed. Ronald Howard and James Matheson. Strategic Decisions Group, 1983: pp 6-15.

- Kahn, Barbara and Robert Meyer. "Consumer Multiattribute Judgments under Attribute-Weight Uncertainty," *Journal of Consumer Research*, 17: 508-522 (March 1991).
- Kalos, Malvin and Paula Whitlock. *Monte Carlo Methods, Volume I: Basics*. New York: John Wiley & Sons, 1986.
- Kassouf, Sheen. *Normative Decision Making*. Englewood Cliffs, NJ: Prentice-Hall Inc., 1970.
- Kaufman, Arnold. *The Science of Decision-Making, an introduction to praxeology*. Toronto, Canada: McGraw-Hill Book Company, 1968.
- Keeney, Ralph L. *Value-Focused Thinking, A Path to Creative Decisionmaking*. Cambridge MA: Harvard University Press, 1992.
- Keeney, Ralph. "Common Mistakes in Making Value Trade-offs," *Operations Research*, 50 (6): 935-945 (November-December 2002).
- Keeney, Ralph and Timothy McDaniels. "Identifying and Structuring Values to Guide Integrated Resource Planning at BC Gas," *Operations Research*, 47 (5): 651-662 (September-October 1999).
- Keeney, Ralph and Howard Raiffa. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Cambridge MA: Cambridge University Press, 1976.
- Kelly, M. and M.C. Thorne. "An approach to multi-attribute utility analysis under parametric uncertainty," *annals of Nuclear Energy*, 28: 875-893 (2001).
- Kim, Soung Hie and Chang Hee Han. "Establishing dominance between alternatives with incomplete information in a hierarchically structured attribute tree," *European Journal of Operational Research*, 122: 79-90 (2000).
- Kirkwood, Craig. *Strategic Decision Making, Multiobjective Decision Analysis with Spreadsheets*. Boston MA: Duxbury Press, 1997.
- Langewisch, Andrew and Fred Choobineh. "Stochastic dominance tests for ranking alternatives under ambiguity," *European Journal of Operational Research*, 95: 139-154 (1996).
- Lavalle, Irving H. *Fundamentals of Decision Analysis*. New York: Holt, Rinehart and Watson, 1978.

- Lavelle, Jerome P., James R. Wilson, Harvey J. Gold, and John R. Canada. "A Method for the Incorporation of Parametric Uncertainty in the Weighted Evaluation Multi-attribute Decision Analysis Model," *Computers & Industrial Engineering*, 32 (4): 769-786 (1997).
- Lee, Kyung Sang, Kyung Sam Park, Yun Seong Eum, and Kwangtae Park. "Extended methods for identifying dominance and potential optimality in multi-criteria analysis with imprecise information," *European Journal of Operational Research*, 134: 557-563 (2001).
- Levary, Reuven and Ke Wan. "A simulation approach for handling uncertainty in the analytic hierarchy process," *European Journal of Operational Research*, 106: 116-122 (1998).
- Moshkovich, Helen, Alexander Mechitov and David Olson. "Ordinal judgments in multiattribute decision analysis," *European Journal of Operational Research*, 137: 625-641 (2002).
- Nauta, Maarten. "Separation of uncertainty and variability in quantitative microbial risk assessment models," *International Journal of Food Microbiology*, 57: 9-18 (2000).
- Olson, D.L., A.I. Mechitov and H. Moshkovich. "Comparison of MCDA Paradigms," *Advances in Decision Analysis*. Edited by: Nadine Meskens and Marc Roubens. Boston MA: Kluwer Academic Publishers, 1999.
- Poyhonen, Mari. *On Attribute Weighting in Value Trees*. Helsinki University of Technology, Systems Analysis Laboratory, Research Reports: Dissertation, March 1998.
- Poyhonen, Mari and Raimo Hamalainen. "On the convergence of multiattribute weighting methods," *European Journal of Operational Research*, 129: 569-585 (2001).
- Raiffa, Howard. "Decision Analysis: A personal account of how it got started and evolved," *Operations Research*, 50 (1): 179-185 (January-February 2002)
- Rothwell, J. Dan. *In Mixed Company: Small Group Communication*. Philadelphia PA: Harcourt Brace Jovanovich College Publishers, 1992.
- Shepetukha, Y. and D. L. Olson. "Comparative Analysis of Multiattribute Techniques Based on Cardinal and Ordinal Inputs," *Mathematical and Computer Modelling*, 34: 229-241 (2001).

Spetzler, Carl and Stael von Holstein, C.S.S. "Probability encoding in Decision Analysis." Presented at the ORSA-TIMS-AIEE Joint National Meeting: Atlantic City NJ, 8-10 November 1972, 603-625.

Stewart, Theodore. "Simplified Approaches for Multicriteria Decision Making under Uncertainty," *Journal of Multi-Criteria Decision Making*, 4: 246-258 (1995).

Wallace, Stein. "Decision Making Under Uncertainty: Is Sensitivity Analysis of Any Use?" *Operations Research*, 48 (1): 20-25 (January-February 2000).

Watson, Stephen R. and Dennis Buede. *Decision Synthesis: The principles and practice of decision analysis*. Cambridge: Cambridge University Press, 1987.

Winston, Wayne L. *Operations Research: Applications and Algorithms, Third Edition*. Belmont CA: Duxbury Press, 1994.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 074-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>					
1. REPORT DATE (DD-MM-YYYY) 23-03-2004		2. REPORT TYPE Master's Thesis		3. DATES COVERED (From – To) Jun 2001 – Jun 2004	
4. TITLE AND SUBTITLE VALUE-FOCUSED THINKING IN THE PRESENCE OF WEIGHT AMBIGUITY: A SOLUTION TECHNIQUE USING MONTE CARLO SIMULATION				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Porter, Timothy, W., Captain, USAF				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 P Street, Building 640 WPAFB OH 45433-8865				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GOR/ENS/04-15	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <p>When a Decision Maker is asked to provide his or her preferences, the response represents a snapshot in time. While their preference structure elicited at any given moment may represent their revealed preferences at that point in time, it may change over time. These changing preferences over time represent ambiguity in the decision maker's preferences. Other sources of ambiguity may exist.</p> <p>One weakness of many decision analysis techniques today is the inability to incorporate ambiguity into the basic decision model. The existence of the problem has been known and commented on for many years. This research addresses that problem. It begins with the basic approach and methodology developed by Ralph Keeney, Value-Focused Thinking (VFT). This methodology is then expanded to allow decision makers to specify not just constant weights to demonstrate their preferences, but an entire distribution. These distributions are then incorporated with the value of the attributes and the whole is simulated using Monte Carlo Simulation provided by Crystal Ball.</p> <p>The result of incorporating these weight distributions into the model, is an empirical distribution for the value of an alternative. The alternative distributions can be compared in a number of ways to provide insight to the decision maker.</p>					
15. SUBJECT TERMS Decision Making, Decision Theory, Uncertainty, Monte Carlo Method, Value-Focused Thinking, Preferences, Ambiguity					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)
U	U	U	UU	99	Stephen P. Chambal, Maj, USAF (ENS) (937) 255-3636, ext 4894 (Kathy McBride)